



15 APRIL 1999 PCT/GB 99 / 01159

The
Patent
Office

91643783



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport

South Wales
RECD 2 JUN 1999

WIPO PCT

6B99 / 1159

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

**PRIORITY
DOCUMENT**

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

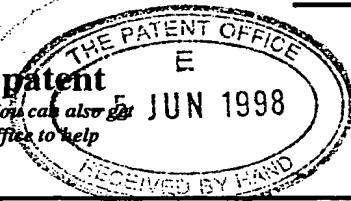
Signed

Dated 11/5/99

This Page Blank (uspto)

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The
Patent
Office

9
08 JUN 98 E366165-1 D03052 The Patent Office
P01/7700 25.00 - 9812162.7 Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

A25641

2. Patent application number
(The Patent Office will fill in this part)

05 JUN 1998

9812162.7

3. Full name, address and postcode of the or of each applicant (underline all surnames)

BRITISH TELECOMMUNICATIONS public limited company
81 NEWGATE STREET
LONDON, EC1A 7AJ, England
Registered in England: 1800000

Patents ADP number (if you know it)

1867002

If the applicant is a corporate body, give the country/state of its incorporation

UNITED KINGDOM

4. Title of the invention

COMMUNICATIONS NETWORK

5. Name of your agent (if you have one)

Timothy Guy Edwin LIDBETTER

"Address for Service" in the United Kingdom
to which all correspondence should be sent (including
the postcode)

BT GROUP LEGAL SERVICES
INTELLECTUAL PROPERTY DEPARTMENT
HOLBORN CENTRE
120 HOLBORN
LONDON, EC1N 2TE

Patents ADP number (if you know it)

1867001

7107774001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country	Priority application number (if you know it)	Date of filing (day / month / year)
---------	---	--

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application	Date of filing (day/month/year)
-------------------------------	------------------------------------

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

YES

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.

See note (d))

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form

Description	44
Claim(s)	7
Abstract	1
Drawing(s)	18

10. If you are also filing any of the following, state how many against each item

Priority Documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature(s)

Date:

5 JUNE 1998

Timothy Guy Edwin LIDBETTER, Authorised Signatory

12. Name and daytime telephone number of person to contact in the United Kingdom

Rohini R Ranjikumar 0171 492 8146

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- a) If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- b) Write your answers in capital letters using black ink or you may type them.
- c) If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- d) If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- e) Once you have filled in the form you must remember to sign and date it.
- f) For details of the fee and ways to pay please contact the Patent Office.

OPTICAL COMMUNICATIONS NETWORK

The present invention relates to an optical communications network, and in particular to the regeneration of optical packets carried on such a network.

It is well known that optical fibre has a huge potential information-carrying capacity. For example, by utilising the entire gain bandwidth of erbium-doped optical amplifiers, a single fibre could carry more than 2 Tbit/s. However in the majority of telecommunications systems in commercial use currently, the information is carried over fibre in the form of an optical signal at a single wavelength. The data transmission bandwidth of the fibre is therefore limited by the electrical bandwidth of the transmitter and receiver, and this means that only a tiny fraction (a maximum of about 1%) of the potential bandwidth-carrying capacity of the fibre is being usefully exploited. There is therefore much interest currently in developing methods for increasing the transmission rate for point-to-point fibre links. One method is wavelength-division multiplexing (WDM), in which several data channels, at different wavelengths, are carried simultaneously on the same fibre. An alternative method for increasing the rate of information that can be carried on fibre is to use optical time-division multiplexing (OTDM) in which several data channels are multiplexed in the form of bit-interleaved return-to-zero (RZ) optical pulse trains.

The WDM approach to photonic networking has some very attractive advantages: in addition to the relative simplicity and commercial availability of the devices needed, WDM networks can be created in a wide variety of architectures with great flexibility (the main restriction being merely that any pair of photonic transmission paths cannot use the same wavelength on a shared fibre link). An advantage of WDM networks is that they can, in principle, support 'signal transparency', i.e. data signals can be carried using any modulation format. However, this implies that, in effect, WDM photonic networks are based on 'analogue' transmission. As a result it is not possible for digital signal regeneration techniques in the optical domain, to be used. The inability to perform signal regeneration in the optical domain leads to practical scaling limitations for WDM networks due to noise accumulation from optical amplifiers, crosstalk and nonlinearity. These factors restrict the number of network switching nodes through which signals can pass without fatal degradation. Currently, in reported laboratory experiments the maximum number of WDM switching nodes through

which a signal can pass without regeneration is limited to around 10, which is a significant restriction in architecture and scalability. A feasible, though costly, solution currently being advocated by some equipment vendors is to sacrifice transparency, standardise the transmission format, and regenerate each 5 wavelength channel individually at the outputs of WDM cross-connects. In effect, this is a hybrid arrangement using analogue switching together with channel-by-channel digital regeneration.

In the OTDM approach to photonic networking, the signals are carried in 'digital' format in the form of RZ optical pulses, allowing the use of digital signal 10 regeneration techniques in the optical domain such as 3R (Re-amplify, Re-time and Re-shape) regeneration [Lucek J K and Smith K, Optics Letters, 18, 1226-28 (1993)] or soliton-control techniques [Ellis A D, Widdowson T, Electronics Letters, 31, 1171-72 (1995)]. These techniques can maintain the integrity of the signals as they pass through a very large number of nodes. For example, Ellis and Widdowson [Ellis A 15 D, Widdowson T, Electronics Letters, 31, 1171-72 (1995)] have made a laboratory demonstration of error-free transmission of signals through an OTDM network consisting of 690 nodes in concatenation. Despite this impressive potential for scalability, however, the OTDM approach to photonic networking suffers from severe restrictions in the network architecture that can be used. This results from 20 the need to maintain proper bit-level synchronism between all the signal sources, demultiplexers and channel add/drop multiplexers throughout the network. In complex architectures, such as one involving the merging of signal streams emanating from several widely-separated sources, fluctuations in the arrival time 25 of signals (due to environmental effects acting on the fibres such as temperature change and mechanical strain) cannot be adequately controlled or compensated in a continuous uninterrupted fashion. This is because of the restricted range of variable optical delay lines, the limited frequency response of control systems due to the physical time of flight of signals over extended distances, and also insufficient degrees of freedom.

30 According to a first aspect of the present invention, there is provided a method of operating a node in an optical communications network including

- a) receiving at the node an optical packet; and

b) generating from the said optical packet received at the said node a regenerated optical packet having a phase determined by a local clock source and independent of the phase of the said packet received at the node.

Throughout this document, the term 'packet' is used to mean a fixed-length or variable-length string of bits which may be routed through a network in a variety of different ways, including self-routing, store-and-forward packet routing, scheduled switched routing and circuit switching

The present invention provides a new approach to operating an optical communications network, which for the first time makes it possible for an optical packet network to be scaled almost without limit. This is achieved by carrying out regeneration of optical packets in a manner which is asynchronous at the bit-level relative to the packet source. This then allows digital signal regeneration to be carried out, facilitating the transmission of packets over large distances, without the extent and architecture of the network being constrained by the need to transmit a global bit-level timing signal.

Preferably the step of generating a regenerated optical packet includes gating, using the received optical packet, an optical clock signal from the local clock source. Preferably this includes passing the optical clock signal through each of a plurality of gate means;

20 applying the received optical packet as a control signal to each of the plurality of gate means with different delays of a fraction of a bit period relative to the optical clock signal input to the gate means; and

25 selecting the output of one of the plurality of gate means to provide the regenerated optical packet. Preferably the difference in delays is equal to T/k where T is the bit period and k is the number of optical gate means. Preferably the width W of the gate window is not less than T/k and not more than T , where T is the bit period and k is the number of optical gates.

30 Alternatively, as further described below, a single gate means may be used in conjunction with means to shift the phase of the incoming packet to match that of the local free-running optical clock.

It is particularly advantageous to use a free-running local optical clock which is gated by the incoming packet. The use of a free-running optical clock, as opposed to one which is phase-locked to an incoming signal, allows the source to be a high quality device such as a passively mode-locked laser. A further

advantage is that the clock may be used to control further optical processing devices at the local node in order, for example, to process header data carried with the packet.

It is found to be particularly effective to use a plurality of optical gates 5 which are gated by the optical packet with different respective delays. The delay in question is that of the optical packet *relative* to the signal from the optical clock source. In practice the different relative delays may be achieved by applying different delays to the optical clock signal, or applying different delays to the optical packet. By using a number of gates in this way, and selecting one of their 10 outputs, it is possible to recover an appropriately regenerated optical signal whatever the phase of the incoming optical packet. The optical gate means may comprise a number of distinct physical devices, such as those described in further detail below. Alternatively the plurality of optical gate means might comprise a single device arranged to gate a plurality of distinct optical signals distinguished, 15 e.g. by their polarisation or wavelength, and references in the description and claims to a number of gate means are to be construed accordingly.

Preferably the method includes making a measurement of a parameter of an optical signal output from the gate means, and selecting the output of one of the plurality of gates to provide the regenerated optical packet depending on the 20 results of the said measurement. The parameter may be a measure of the total energy of the output signal, and the selection may be made by comparing the energies of the signals from different optical gate means. Other parameters may also be used to make the selection. For example, the bit error level may be measured, and the output selected which has the minimum bit error level.

25 According to a second aspect of the present invention, there is provided a method of operating a communications network comprising a plurality of nodes interconnected by an optical transmission medium, the method including:

transmitting an optical packet and at a network node, receiving the said packet and generating from the said packet a regenerated optical packet having a 30 phase determined by a local optical clock source and independent of the phase of the said packet received at the network node.

According to a third aspect of the present invention, there is provided a regenerator for optical packets comprising:

an input for receiving an optical packet;

a local optical clock source comprising a free-running local oscillator; means for generating from the optical packet and from a clock signal from the local optical clock source a regenerated optical packet independent in phase from the optical packet received at the input.

5 Systems embodying the present invention will now be described in further detail, by way of example only, and will be contrasted with the prior art, with reference to the accompanying drawings, in which:

Figure 1: shows the basic concept for 3R regeneration of a digital data stream consisting of a RZ pulse train encoded by on-off modulation

10 Figure 2: shows the method of 3R regeneration used for OTDM systems in which the bit stream is continuous

Figure 3: shows a method of partial regeneration of optical packets

Figure 4: a simplified outline diagram showing a generator of optical packets and a bit-asynchronous packet regenerator

15 Figure 5: a dual-gate bit-asynchronous packet regenerator

Figure 6: a sequence of timing diagrams that illustrate the operation of the dual-gate regenerator

Figure 7: diagram showing the probability density function of the arrival time of the i th data bit at the gates, relative to the clock signal input to gate G1

20 Figure 8: a plot of the expression (12) against the phase angle θ , taking the values $\sigma = 0.018T$ and $W = 0.75T$

Figure 9: shows a plot of the bit-error probability B against the phase angle θ , calculated according to (12), using an ideal mechanism to select the output gate i , and taking $W/T = 0.75$ and $\sigma/T = 0.018$ and 0.03

25 Figure 10 shows a plot of the maximum bit-error probability B according to (12), plotted versus the bit-arrival jitter σ , and taking $W/T = 0.75$.

Figure 11: is a plot of the expected bit-error probability for a given packet, per regenerator passed, according to (15), plotted against the rms jitter σ , with $W/T = 0.65, 0.75, 0.85$.

30 Figure 12: shows a plot of the bit-error probability B against the phase angle θ , calculated according to (18), taking $W/T = 0.75$, $\sigma/T = 0.018$ and $\sigma_c/T = 0.001$

Figure 13 shows a plot of the maximum value of the bit-error probability B for the phase angle θ anywhere in the range $0 \leq \theta < 2\pi$, calculated according to (18), plotted versus σ_c/T , with $W/T = 0.75$ and $\sigma/T = 0.001, 0.018$ and 0.03 .

Figure 14: Diagram of a bit-asynchronous quad-gate packet regenerator

Figure 15: plot of bit-error probability according to (19) plotted against the phase angle θ , taking $\sigma = 0.036T$ and $W = 0.75T$

Figure 16: a plot of the bit-error probability B against the phase angle θ ,
5 calculated according to (12), using an ideal mechanism to select the output gate / (such as the strategy (20) executed by ideal circuits without systematic or random errors) and taking $\sigma = 0.036T$ and $W = 0.75T$

Figure 17: Plot of the bit-error probability B against the phase angle θ , calculated according to (21), and assuming that $W/T = 0.75$, $\sigma/T = 0.036$ and
10 $\sigma_c/T = 0.001$ and 0.052 ."

Figure 18: Plot of the maximum value of the bit-error probability B for any phase angle θ in the range $0 \leq \theta < 2\pi$, calculated according to (21), plotted versus W/T , with $\sigma/T = 0.036$, and $\sigma_c/T = 0.001$ and 0.05 ."

Figure 19: show values for the probability that a packet suffers 'slippage' (i.e. a
15 time displacement greater than a specified acceptable limit $L7$), according to (22), plotted versus N , the number of regenerators passed, for various values of $L7$.

Figure 20: shows a version of the two-gate regenerator which uses only one optical gating device to gate simultaneously two independent clock signals which are distinguishable by their different states of polarisation.

20 Figure 21: Outline diagram of a bit-asynchronous packet regenerator using a single gate to modulate the output of the local source.

Figure 22: Example embodiment of the bit-asynchronous packet regenerator using a single gate to modulate the output of the local source.

Figure 23: Example layout of a portion of a network comprising switching nodes
25 containing routing switches (RS) and bit-asynchronous regenerators (AR), and links between the switching nodes containing bit-synchronous regenerators (SR).

Figure 24: Outline diagram of an alternative arrangement of the bit-asynchronous packet regenerator using a single gate to modulate the output of the local source.

30 Figure 25: A further example embodiment of the bit-asynchronous packet regenerator using a single gate to modulate the output of the local source

Figure 26 shows a schematic of an optical packet network.

In an ultrafast optical packet network, information is transported across a network in the form of fixed-length bursts (i.e. cells or fixed-length packets) of RZ

(return-to-zero) optical pulses that are encoded with payload data and control information (such as the address of the packet destination). The network may constitute, for example, the core network of a national data/telephony network, or a local area network connecting a number of computing systems, or part of the

5 fabric of a communications switch, or may provide the connection between processors in a multi-processor computer. Examples of suitable topologies for such a network are described in the present applicant's co-pending European patent application 97307224.2, the contents of which are incorporated herein by reference. The use of logical (header), rather than physical, addressing facilitates

10 massive scalability. Every transmission path in the network carries a continuous sequence of time slots, synchronised to a packet-level global clock, and each time slot accommodates at most one packet and an appropriate time guard band. This time guard band allows the transmission path to be switched for packet-by-packet routing and also allows continuous and endless synchronisation of the packet

15 streams and network nodes to the packet-level global clock. Thus the packet time slots are synchronised throughout the network. However, a crucial aspect of our approach is that synchronism between packets at the (picosecond) bit-level is not required. Therefore the portion of the time slot that accommodates the packet can be conveniently defined to be several bit periods longer than the packet duration.

20 This permits a certain amount of slop in the positioning of the packet within the time slot - in other words, the positioning of the packet within the time slot is not made with bit-level precision, and generally successive packets are not bit-synchronous.

As is further described below, the use of a bit-asynchronous digital optical

25 packet regenerator, makes it possible to scale upwards such a network almost without limit.

Aspects of the invention are described, by way of example only, in the following sections.

JOURNAL OF LIGHTWAVE TECHNOLOGY

Special Issue on Photonic Packet Switching Systems, Technologies and Techniques

Asynchronous Digital Optical Regeneration and Networks (Invited Paper)

David Cotter and Andrew D Ellis

BT Laboratories, Martlesham Heath, Ipswich, Suffolk IP5 3RE, UK

Abstract—This paper outlines the concept of the asynchronous digital optical network, which is aimed to combine the advantages, and overcome the drawbacks, of existing approaches to photonic networking. The network is based on digital optical transmission and processing, which offers the prospect of almost infinite scalability. Optical transmission throughout the network is in burst mode, using a standard digital optical signal format. The main defining feature of the asynchronous network is that the network nodes (in which the optical processing and routing take place) do not share global synchronisation at the bit level. Instead, the nodes each operate with independent bit-level clocks. This approach removes the architectural constraints found in synchronous networks. An essential component of the network is the asynchronous digital optical regenerator. Various design and performance aspects of this new type of regenerator are analysed.

Index Terms—Photonic network, digital optical network, asynchronous network, optical regenerator, burst-mode transmission, photonic packet network

I. INTRODUCTION

Techniques such as wavelength-division multiplexing (WDM) and optical time-division multiplexing (OTDM) are highly effective in increasing the bandwidth of point-to-point optical transmission systems, but tend to exacerbate the significance of the ‘electronic bottleneck’ in communications networks—i.e. the limited capacity of electronic switching and processing systems attached to the ends of the optical links. To avoid these bottlenecks, there is much interest in developing photonic networks, in

which information is carried from its source to destination in optical form without the need for optoelectronic conversion or electronic processing at intermediate nodes. This could allow much greater overall network capacities to be achieved, with smaller and less costly electronic switches and processors, greater flexibility and simpler management. Up to now, photonic networks based on carrier (i.e. wavelength) routing and transparent optical cross-connects have received by far the most attention from researchers. The devices needed for combining and splitting optical channels defined by wavelength are already well developed, and early trials and deployments of wavelength-routed networks are currently in progress in various countries. On the other hand, optical channel add/drop multiplexers operating at the bit level (such as in OTDM) are rare research devices available in only a few laboratories. Although there have been many demonstrations of OTDM point-to-point transmission, so far there have been only very few demonstrations of OTDM networking [1-6].

The wavelength-routing approach to photonic networks has some very attractive advantages: in addition to the relative simplicity and commercial availability of the components needed, a great variety of network architectures can be used (the main restriction being merely that any two transmission paths cannot use the same wavelength on a shared fibre link). A much-vaunted advantage of wavelength-routed networks is that they can support 'signal transparency', i.e. data signals can be carried using any modulation format. However, this implies 'analogue' transmission, which precludes digital techniques such as 3R regeneration in the optical domain. The inability to perform analogue optical regeneration leads to practical scaling limitations due to noise accumulation from optical amplifiers, dispersion, crosstalk, nonlinearity and other imperfections in transmission and cross-connects. All of these restrict the number of network nodes through which signals can pass without fatal degradation. Currently, in reported laboratory experiments the maximum number of wavelength-routing nodes through which a signal can pass without regeneration is limited to around 10, which is a significant restriction in both scalability and operational flexibility [7]. A feasible, though costly, alternative currently advocated by some equipment vendors is to retain wavelength-routing, but sacrifice transparency, standardise the transmission format, and regenerate every wavelength channel individually at each node [8]. In effect, this is a hybrid arrangement using

analogue cross-connects together with channel-by-channel digital regeneration. The practical scalability of this solution is uncertain, as also are the advantages to be gained from retaining analogue cross-connects.

In the 'digital' approach to photonic networking (as in OTDM networks) the signals are carried in a standard binary format, allowing ultrafast digital signal processing techniques in the optical domain (e.g. 3R regeneration [9-11], soliton-control techniques [2], and header address recognition [12]). The use of digital optical regeneration can maintain the integrity of ultra-high bandwidth signals as they pass through a very large number of network nodes. For example, error-free transmission of digital optical signals, without optoelectronic conversion, through 690 nodes in an OTDM network has been demonstrated [2]. However, OTDM in its traditional form requires bit-level synchronism to be maintained between all the signal sources, demultiplexers and channel add/drop multiplexers throughout the network, and this imposes severe restrictions on the architecture that can be used. For example, the 690-node network [2] consisted simply of a linear chain of add/drop multiplexers with locally synchronised sources. A more complex and practical architecture for the wide area, such as a mesh network, requires signal streams emanating from many widely-separated sources to be merged at several positions. In this case, variations in the arrival time of signals (due to phase fluctuations in the sources as well as environmental effects acting on the transmission fibre) cannot be adequately controlled or compensated in an uninterrupted fashion. This is due to the lack of variable buffer stores in the optical domain, the restricted physical range of variable optical delay lines, and the limited frequency response of control systems for source synchronisation (because of the fundamental time-of-flight loop delay over extended distances) [6]. Consequently there are insufficient degrees of freedom to allow anything other than the most simple architectures (bus or ring).

Thus the advantages and disadvantages of the existing approaches to photonic networking appear to be complementary. The 'transparent' wavelength-routing approach to networking is asynchronous (bit-level synchronisation between different network nodes and channels is not required) making for network flexibility, but appears not to be readily scalable. On the other hand, the digital approach to photonic networking allows almost infinite scalability both in terms of distance and

the numbers of nodes, but the network architecture is severely limited if bit-synchronism must be maintained across the network (as in OTDM). This paper proposes an alternative approach, which we term the 'asynchronous digital photonic network'. The aim is to combine the advantages of high-speed digital optical processing with the architectural flexibility of a network which does not require bit-level synchronisation between the nodes.

Section II describes the concept of asynchronous digital photonic networks. Section III introduces an essential component of the network: the asynchronous digital optical regenerator. Aspects of the design and predicted performance of such regenerators are presented. A simple example of a network implementation is given in Section IV.

II. ASYNCHRONOUS DIGITAL PHOTONIC NETWORKS

An asynchronous digital photonic network is one in which there is no global synchronisation at the bit level. In each node the bit-level clock (used to drive the subsystems, such as sources, routers and demultiplexers within the node) operates at the standardised nominal bit rate, but there is no attempt made to achieve phase synchronisation between bit-level clocks in different nodes. The physical optical transmission across the network occurs in burst mode, where a 'burst' is a finite contiguous string of bits in a standard signal format, which is assumed here to be return-to-zero (RZ) optical pulses at the nominal bit rate. The burst may represent an individual packet or cell (with header and payload), or could contain data belonging to several multiplexed channels or virtual circuits. Routing across the network could be achieved by any appropriate method, such as circuit switching, packet or cell switching, self-routing, just-in-time switching, scheduled routing, etc [13].

Throughout the network, time is divided into slots. Although the nodes are asynchronous at the bit level, the network operates in a synchronous fashion at the slot level to avoid contention at the routing nodes [14]. The time slots are of fixed duration, and all of the nodes and links operate to a global clock at the time-slot rate. The slots are dimensioned to contain at most one burst, together with appropriate time guard bands to maintain proper separation of bursts and to allow time for the

operation of routing switches (Fig. 1). It is envisaged that the burst length might be fixed at a value between perhaps 1,000 and 100,000 bits at a peak bit rate of 100 Gbit/s or greater, and consequently the time-slot duration might be fixed at between 10 ns and 1 μ s. However the choice of these parameters and the method of routing the bursts will depend on the types of traffic carried and many other issues, and are not considered here.

The slot-level synchronisation ensures that when a burst is introduced onto the network it is inserted into the appropriate part of a vacant time slot. However the global time-slot clock is not defined with bit-level (picosecond) precision, and there is no correlation between the time-slot clock and bit-level clocks. Consequently the position of a burst relative to the time slot boundaries could vary by as much as tens of bit periods, and due allowance for this should be made in dimensioning the time slot (Fig. 1).

By dispensing with global bit-level synchronisation, the asynchronous digital photonic network is free of the architectural constraints found in synchronous networks. Although the need for global synchronisation at the slot level remains, this is relatively easily achieved when using burst-mode transmission [15]. The lengths of the links between nodes are controlled so that the time-slot boundaries at each input to a node are aligned to the time-slot clock. This can be achieved, for example, by using a switched optical delay unit [15] which is operated during the guard band between bursts.

Since the bursts arriving at a network node are generally not in synchronism with the local bit-level clock, this increases the difficulty of bit-level processing such as regeneration, routing, demultiplexing, etc, because phase acquisition is required within each time slot. Consequently, a key sub-system of the nodes in the asynchronous digital photonic network is the asynchronous digital burst regenerator, described below. As will become clear later, in addition to providing a simpler technique for burst phase acquisition and regeneration, this type of regenerator is capable of re-synchronising all the incoming traffic to the local clock, thus greatly simplifying subsequent bit-level processing at the node.

III. ASYNCHRONOUS DIGITAL BURST REGENERATOR

A. Principle of Operation

Figure 2 shows the well-known concept for optical regeneration of a RZ digital data stream [9-11]. The incoming data bits from a distant source are used to modulate a continuous train of high-quality pulses produced by a synchronous local source, thus regenerating the original data. Each 'mark' in the incoming data causes the gate to switch to transmission mode for a fixed time (the gate window), allowing a single pulse from the local source to pass through. In this way the regenerated bits have substantially the same pulse shape, spectral quality, amplitude and timing stability as the local source. Moreover the regenerator can tolerate a degree of jitter in the arrival time of the data bits, determined by the gate window width. A key design problem in the regenerator is to ensure proper synchronisation between the local pulse source and the incoming data bits. In a conventional (synchronous) regenerator this is achieved by forcing the local source into synchronism with the incoming bit stream, usually by some form of clock recovery. However in the asynchronous regenerator that we introduce here, the local pulse source is free-running (with a repetition frequency close to the nominal bit rate) and no clock recovery is used. Instead of forcing the local pulse source into synchronism with the incoming data as in the synchronous regenerator, the asynchronous regenerator allows for time-varying phase differences between the incoming data bursts and the local pulse source. After regeneration, all data bursts, regardless of origin, take up the frequency and phase of the independent local source.

The bit rate of an incoming data burst at the asynchronous regenerator (f_B) and the frequency of the local clock (f_C) are nominally the same as the standard bit rate $1/T$, where T is the nominal bit period, but in general they differ slightly, i.e.

$f_B \approx f_C \approx 1/T$. Therefore the instantaneous phase difference between any incoming burst and the local clock is random. However, provided the burst is sufficiently short, this phase difference will not change significantly during the burst. To quantify this statement, consider the phase difference between the local clock pulses and the incoming data bits at the inputs to the gate in the regenerator (Fig. 2), as a function of time: $\theta(t) = 2\pi\Delta f t + \theta_0$, where $\Delta f = f_C - f_B$ and θ_0 is a constant offset. The

change in $\theta(t)$ over the full duration of a burst of length n bits is $\Delta\theta = 2\pi\Delta f(n-1)T$. This corresponds to a time slippage Δt in the local clock over the duration of the burst,

$$\Delta t = (\Delta\theta/2\pi)T = \Delta f(n-1)T^2, \quad (1)$$

which can be considered negligible if $|\Delta f| \ll T$. Figure 3 shows a plot of $|\Delta t|/T$ versus the burst length n , for various values of the nominal bit rate, assuming that $|\Delta f| = 50$ kHz. This value of $|\Delta f|$ can be guaranteed for free-running network nodes [16]. At 100 Gbit/s, for example, and for a burst length $n = 10,000$ bits, $|\Delta t| = 0.005T = 50$ fs. This shows that with realistic parameters, the time slippage over the duration of a burst can be negligibly small.

Therefore, we can make the simplifying assumption that the phase difference θ between the incoming data and local clock is constant over the duration of the burst. However θ is unknown, and the asynchronous regenerator should be designed to regenerate correctly each burst for any value of θ in $[0, 2\pi]$. This can be achieved in different ways. One approach, which is particularly simple and robust, is to use a regenerator with multiple gates, and this is described in IIIB-D.

This leads to a hypothesis where each node in the network contains a single, independent clock source. In general each incoming burst, on each input fibre, arrives with an arbitrary phase determined by the unknown phase of the source clock and the network path taken. Asynchronous regenerators, driven by the same local clock, are placed at each of the inputs to the node. The asynchronous regenerator has the property that the regenerated bursts are all in precise synchronism with the local clock. Consequently, the most important advantage of using asynchronous regenerators is that all the bursts arriving on the various incoming links at a network node are brought into synchronism. This allows the nodes to operate without network-wide bit-level synchronisation, avoiding the drawbacks of bit-synchronous networks. A further advantage is that there is now considerable freedom in the choice of optical pulse source to act as the local clock, without regard to ease of synchronisation to an external frequency reference, since this is no longer required. The optical pulse source may be an RF-driven laser, or alternatively a free-running laser may be used (for example, a stable phase-locked erbium fibre ring laser [17]). Similarly, if external

pulse compression is required [18-20] the inevitable phase fluctuations associated with long lengths of fibre may be neglected. Therefore the source may be selected mainly on its suitability for providing the required quality of pulses for regeneration, processing and onward transmission. A further important feature of the asynchronous regenerator is that, because the regenerated bursts are all in precise phase-synchronism with the local independent pulse source, this source can provide a continuous supply of synchronised optical pulses for subsequent bit-level processing in the node (such as 'on-the-fly' error detection, demultiplexing, burst header processing, etc) [21]. Finally, it should be noted that since the phase of each regenerated burst within a node is well defined with reference to the local clock, it is possible (but not necessary) to arrange for all the bursts leaving the node along a particular transmission line to be phase coherent. Thus, if regeneration is required during subsequent transmission on that line, rapid burst-to-burst phase acquisition is not needed for successful clock recovery, and traditional synchronous techniques may be applied.

B. Dual-Gate Regenerator

The asynchronous regenerator must be designed to regenerate correctly each burst regardless of the phase difference θ between the burst and the local clock. This cannot be achieved by providing just one gate (as shown in Fig. 2) because even for an ideal square switching window of width W equal to the bit period T , ambiguity exists at the edge of the switching window. However, two or more gates, each of which can correctly regenerate a burst for only a restricted range of θ , can together span the complete interval $[0, 2\pi]$. By monitoring the gate outputs, the regenerator selects a correctly regenerated output from one of the gates in each time slot.

The simplest form is a regenerator containing only two gates. As will become apparent, this has practical deficiencies. Nevertheless it serves to illustrate the design problems, leading to the more robust quad-gate version described in IIIC.

1) Principle of Operation: Figure 4 shows a dual-gate burst regenerator. The data bits in the incoming burst are used to control simultaneously the operation of two gates, G1 and G2, which are assumed here to be non-inverting and have gate windows of equal width W . The output from the local clock is applied to the inputs of the two

gates, one of these inputs being delayed relative to the other by an amount $T/2$ (modulo T). Since the phase θ of the local clock pulses relative to the incoming data bits has an arbitrary and unknown value in the interval $[0, 2\pi]$, it is necessary that the gate window widths are chosen so that, whatever the value of θ , the clock pulses will be correctly modulated by at least one of the two gates. If the clock pulses are sufficiently narrow that they can be considered to be impulses, the window width W should be in the range $T/2 < W < T$. The lower limit ensures that at least one clock pulse will be modulated whatever the value of θ , whilst the upper limit comes from the requirement that no more than one clock pulse may pass through a gate during the window.

The optical gates may be implemented in different ways. For ultrafast operation, the gates could be non-linear optical devices such as the fibre loop mirror [22–24], in which case the window width W is defined by selecting suitable fibre lengths, dispersion and birefringence. Alternatively gates based on the nonlinearity in semiconductor optical amplifiers could be used [25, 26], where the gate width is largely determined by the geometry of an interferometer. Dual ‘gates’ can be provided by a single gating device by switching simultaneously two optical clock signals which are separable, for example by having orthogonal states of polarisation or different directions of propagation [27].

Figure 5 shows a sequence of timing diagrams that illustrate the operation of the dual-gate regenerator. The burst data bits (an example sequence 11101 is shown) arrive at the control ports of the gates G1 and G2, and each ‘mark’ causes the gates to switch for a time W . The diagrams (i–iv) illustrate various values of θ , the phase of the local clock relative to the incoming burst data bits. In case (ii) where $W/T - 1/2 < \theta/2\pi < 1/2$, or case (iv) where $W/T < \theta/2\pi < 1$, the phase is such that one or other of the gates correctly modulates the clock pulse train so as to produce a regenerated copy of the data bits at the output. In case (i) where $0 < \theta/2\pi < W/T - 1/2$, or case (iii) where $1/2 < \theta/2\pi < W/T$, the outputs from both gates G1 and G2 are correctly modulated simultaneously.

Returning to Fig. 4, the components shown to the right-hand side of the two gates are used to attempt to select in each time slot whichever gate output gives a

regenerated burst with the minimum of bit errors. One technique is to make the selection in each time slot on the basis of a comparison of the total optical energy emerging from each gate, integrated over the duration of the burst. If the phase angle θ is such that the output from a gate consists of correctly modulated clock pulses then the total optical energy measured at the output of the gate, integrated over the burst, will be maximum (in effect, it is a measure of the number of 'marks' appearing in the regenerated data burst). However if θ is such that the clock pulses arrive at the gate at a time outside the gate window, then the energy transmitted by the gate will be small or zero. The circuit shown in Fig. 4 therefore makes and compares these energy measurements, and the result is fed forward to operate an optoelectronic switch S which makes the physical selection of the gate output. The detectors, switch and associated control systems, which operate at the time-slot rate (not the data bit rate), are synchronised to the global time-slot clock, and the selector changes state during the guard bands between bursts. The measurements and control operations may be broken into pipelined stages, each stage requiring a time of less than one time slot. To allow time for these pipelined operations, in Fig. 4 additional optical delays are provided between the outputs of the gates and the selector switch.

The precise optical delays inside the regenerator are arranged so that the relative time delay $T/2$ for the clock pulses at the inputs to the gates is exactly compensated at the output of the regenerator (this is the reason in Fig. 4 for the additional relative delay of $T/2$ in the output from gate G1). Although not strictly necessary for the operation of the regenerator, the advantage of doing this is that at the output all the bursts are in phase synchronism with each other and with the local clock, and as mentioned earlier, this simplifies subsequent digital optical processing. A further important benefit of equalising the path delays in this way is to limit the amount of 'slippage' in a large network, as described in IIIF.

For a high speed (100 Gbit/s) system, there are two main causes of bit errors that may occur in the process of regenerating a burst using the asynchronous regenerator. The first cause is jitter in the arrival time of the incoming burst data bits. Some degree of jitter is inherent in all sources, and it is also well known that additional timing jitter arises during transmission from effects such as spontaneous emission noise, the self-frequency shift due to stimulated Raman amplification, short-range interactions

between solitons, and the complex interplay between these and other processes [28]. If the time of arrival of a data 'mark' at the regenerator fluctuates relative to the local clock pulse train, the time position of the gate window opened by that data bit is shifted accordingly. This increases the probability that a clock pulse may fall outside the gate window and so may fail to be transmitted correctly as a 'mark' in the regenerated burst. The second main cause of bit errors in the asynchronous regenerator is errors in selecting in each time slot the most appropriate gate from which to take the output. The bit-error probability arising from timing jitter is now analysed in detail under the assumption that the gate-selection mechanism is free of systematic and random errors and has infinitesimal resolution. In a succeeding subsection the effect of errors in the operation of a realistic gate selector is considered.

2) Bit-Error Probability at the Output of Each Gate: In IIIA we justified the assumption that the phase difference θ between an incoming data burst and the local clock can be taken to be constant over the duration of the burst. However, as explained above, the arrival time of the data bits in an individual burst are subject to bit-to-bit fluctuations relative to the local clock. Figure 6 shows a probability density function for the arrival time of the i th data bit at the gates, relative to the clock signal input to gate G1 [24]. This relative arrival time is taken to be a normally distributed random variable with standard deviation σ (usually referred to as the rms jitter), and thus σ incorporates jitter in the absolute arrival time of data-bits as well as jitter in the local clock. The mean value of the distribution occurs at a time $\theta T / 2\pi$ before a clock pulse. If the actual arrival time of the i th data bit is within the unshaded region in Fig. 6 defined by $t_i - W < t < t_i$ (where t_i is the time of the i th clock pulse) and if the bit is a 'mark', the gate will allow the i th clock pulse to be transmitted correctly thus regenerating the data bit. If the arrival time of the i th data bit falls outside this unshaded region, then the i th clock pulse is not correctly modulated. The probability that the actual arrival time falls within the unshaded region is given by

$$p(\psi) = \Phi\left(\frac{\psi}{\sigma/T}\right) - \Phi\left(\frac{\psi - W/T}{\sigma/T}\right) \quad (2)$$

with $\psi = \psi_1$, where

$$\psi_1 = \theta/2\pi, \quad (3)$$

and $\Phi(z)$ is the standard normal cumulative distribution function:

$$\Phi(z) = \int_{-\infty}^z \phi(x) dx , \quad (4)$$

$$\phi(z) = \exp(-z^2/2) / \sqrt{2\pi} . \quad (5)$$

If the actual arrival time of the i th data bit is within the heavily shaded region in Fig. 6 defined by $t < t_{i-1} = t_i - T$ and if the bit is a 'mark', the gate will allow the $(i-1)$ th clock pulse to be transmitted, with the possibility of producing an error. The probability that the actual arrival time falls within the heavily shaded region is given by

$$q(\psi) = \Phi\left(\frac{\psi - 1}{\sigma/T}\right) \quad (6)$$

with $\psi = \psi_1$.

Similarly, if the actual arrival time of the i th data bit is within the other heavily shaded region defined by $t > t_{i+1} - W = t_i + T - W$ and if the bit is a 'mark', the gate will allow the $(i+1)$ th clock pulse to be transmitted, again with the possibility of producing an error. The probability that the actual arrival time falls within this region is given by

$$r(\psi) = 1 - \Phi\left(\frac{\psi + 1 - W/T}{\sigma/T}\right) \quad (7)$$

with $\psi = \psi_1$.

At the gate G2, the clock pulses are arranged to arrive later in time, relative to G1, by an amount $T/2$ (modulo T). In this case the expressions for the probabilities p , q and r are given by (2), (6) and (7), but with $\psi = \psi_2$ where

$$\psi_2 = (\theta/2\pi + 1/2) \bmod 1 \quad (8)$$

For both gates, the bit error probability can be calculated by considering the various combinations of values of the j th data bit with its nearest neighbours, as presented in Table 1. It is assumed that each data bit is equally likely to be a 0 or 1, and Table 1 neglects the small reduction in bit-error probability for the end bits, $j = 1$ or n , where there are n bits in the burst. The parameter α takes a value between 0 and 1,

depending on the type of gate and its action in response to two control pulses separated by less than W . By summing the probabilities listed in Table 1, the overall bit-error probability for the output from gate i is:

$$\begin{aligned}\langle e | \theta \rangle &= (1 - p)/2 + (2 - \alpha)p(q + r)/4 - (4 - 3\alpha)pqr/8 \\ &\approx (1 - p)/2 ,\end{aligned}\quad (9)$$

where e is a random variable that takes the value 1 when a bit error occurs and the value 0 otherwise, and $\langle e | \theta \rangle$ denotes the expected value of e conditional on the phase angle θ . The approximate version of (9) is good for $\sigma/T < 0.1$, regardless of the value of α .

Figure 7 is a plot of the error probability $\langle e \rangle$ against the phase angle θ , taking as an example the values $\sigma = 0.018T$ and $W = 0.75T$ for the jitter and gate window width, respectively. It can be seen that, with these parameter values, for any value of θ in $[0, 2\pi]$, the bit-error probability for the output from at least one of the gates is expected to be less than 10^{-12} . Therefore this bit-error probability could be achieved continuously by selecting the output from one of the gates suitably in each time slot. In this work we take 10^{-12} as the target bit-error probability (or the equivalent probability of $n \times 10^{-12}$ that a burst of n bits will contain an error).

3) *Ideal Gate Selection:* We now assume that the output gate selection is based on a comparison, in each time slot, of the energies of the burst emerging from each gate, as described in IV A. As noted earlier, the total optical energy measured at the output of a gate, integrated over the duration of the burst, is in effect a measure of the number of 'marks' appearing in the regenerated data burst. Similar to the calculation of the bit error probability given above, the expected value of the optical energy measured at the output of the i th gate can be calculated by considering the various combinations of values of the j th data bit with its nearest neighbours, as presented in Table 2. Again it is assumed that each data bit is equally likely to be a 0 or 1. The expected value of the optical energy measured at the output of the i th gate, $\langle E_i \rangle$, integrated over the duration of the burst, is found by summing the probabilities in the table:

$$\begin{aligned}\langle E_i | \theta \rangle &= (p + q + r) / 2 - (2 - \alpha)(pq + pr + qr) / 4 + 3(1 - \alpha)pqr / 8 \\ &\approx (p + q + r) / 2.\end{aligned}\quad (10)$$

The approximation is good for $\sigma/T < 0.1$, regardless of the value of α . Here $\langle E_i \rangle$ is normalised to the value nw , where n is the number of bits in the original burst and w is the energy of a pulse representing a 'mark'. Then the strategy for selecting the gate output in each time slot is:

```
if  $\langle E_1 \rangle > \langle E_2 \rangle$  then select gate 1
else select gate 2.
```

Figure 8 shows a plot of the bit-error probability $\langle e | \theta \rangle$ calculated according to (9), based on the strategy (11) under the assumption that the gate-selection mechanism operates perfectly in every time slot. It can be seen that in the case where the jitter is $\sigma = 0.018T$, the maximum value of $\langle e | \theta \rangle$ is expected to be 10^{-12} (as in Fig. 7). The other case shown demonstrates that the predicted maximum bit-error probability increases with rms jitter σ .

4) *Imperfect Gate Selection*: The results presented above assume that the mechanism for selecting the output gate is ideal (i.e. that there is infinitesimal resolution and no systematic or random errors). More realistically, the gate selection will be imperfect, so that the incorrect gate may be selected at certain times leading to an increase in the bit-error probability.

In the dual-gate regenerator shown in Fig. 4, the energies of the regenerated bursts emerging from gate G1 and G2 are measured in each time slot, and the results of these measurements, E_1 and E_2 respectively, are compared and used to control the operation of the selector switch S according to the strategy (11). The energy measurements and comparison are subject to systematic and random errors. Here we assume that systematic errors are negligible, and model the random errors by assuming that, for any given burst with phase angle θ , the measured value $E_1 - E_2$ is a normally-distributed random variable with mean $\langle E_1 | \theta \rangle - \langle E_2 | \theta \rangle$ and variance σ_c^2 . On the basis of these measurements, gate G1 will be selected with the probability:

$$P(G1) = \Phi[(\langle E_1 | \theta \rangle - \langle E_2 | \theta \rangle) / \sigma_c], \quad (12)$$

where $\Phi(z)$ is given by (4). The overall bit-error probability is therefore

$$\langle e | \theta \rangle = P(G1) \langle e | \psi_1 \rangle + [1 - P(G1)] \langle e | \psi_2 \rangle, \quad (13)$$

where $\langle e | \psi_i \rangle$ is given by (9), with ψ_i given by (3) and (8). Figure 9 shows a plot of the bit-error probability $\langle e \rangle$, assuming a small value of 0.1% for the measurement standard deviation σ_c , and taking for example $W = 0.75T$ and $\sigma = 0.018T$. By comparing with the bit-error probability for the dual-gate regenerator with ideal gate selection (Fig. 8), it is seen that the effect of small imperfections in the gate-selection mechanism can cause severe degradation in performance. Indeed, taking the example with the input rms jitter $\sigma = 0.018T$, a value of the measurement standard deviation σ_c as small as 10^{-7} is sufficient to degrade the bit-error probability by 4 orders of magnitude (from 10^{-12} in the ideal case to around 10^{-8}). This is a serious limitation in the practical usefulness of the dual-gate asynchronous regenerator. The quad-gate design described below overcomes this limitation.

C. Quad-Gate Regenerator

1) *Principle of operation:* Figure 10 shows a quad-gate asynchronous burst regenerator. The principle of operation is similar to the dual-gate regenerator, except that the incoming data bits control the operation of four gates simultaneously, and the output from the local clock is applied to the inputs of the four gates with relative delays in steps of $T/4$ (modulo T). Similarly to the dual-gate regenerator, to ensure that the clock pulses will be correctly modulated by at least one of the four gates whatever the value of θ , the gate window width W must lie within the range $T/4 < W < T$. We simplify the analysis by assuming that the incoming data pulses and local clock pulses can be represented by impulses. As in the dual-gate regenerator, the technique for gate selection shown in Fig. 10 is based on a comparison of the optical energy emerging from each gate, integrated over the duration of the burst.

2) *Bit-Error Probability:* The bit-error probability $\langle e \rangle$ for the output from gate i is given by (9) with $\psi = \psi_i$, where

$$\psi_i = [\theta/2\pi + (i-1)/4] \bmod 1, \quad (14)$$

with $i = 1, \dots, 4$. Figure 11 is a plot of $\langle e \rangle$ against the phase angle θ , taking as an example the values $\sigma = 0.036T$ and $W = 0.75T$ for the jitter and gate window width, respectively. It can be seen that, with these parameter values, for any value of θ in $[0, 2\pi]$ the bit-error probability for the output from at least one of the gates is less than 10^{-12} .

An appropriate strategy for selecting the gate output in each time slot is to select the gate whose clock input is 180° out of phase with the clock input to the gate with the lowest-energy output, i.e.:

$$\text{if } \langle E_j \rangle = \min \{ \langle E_i \rangle \} \text{ then select gate } (1 + (j + 1) \bmod 4), \quad (15)$$

where $\{ \langle E_i \rangle \}$ denotes the set $\{ \langle E_1 \rangle, \langle E_2 \rangle, \langle E_3 \rangle, \langle E_4 \rangle \}$. The expected value of the optical energy measured at the output of the i th gate, $\langle E_i \rangle$, is given by (10). Figure 12 shows a plot of the bit-error probability $\langle e \rangle$ against phase angle θ , using the strategy (15) to select the output gate i in an ideal fashion without systematic or random errors ($\sigma_c = 0$). A maximum bit-error probability of 10^{-12} is predicted when the rms jitter $\sigma = 0.036T$. Comparing this with the performance of the dual-gate regenerator (Fig. 9), the quad-gate regenerator can tolerate a higher level of jitter. The reason is that, in the case of the quad-gate regenerator, the clock pulses modulated by the selected gate are positioned further from the edges of the gate window.

As in IVB4, we now consider the effect on the performance of the regenerator when the mechanism for selecting the output gate is imperfect. We assume that systematic errors are negligible, and consider the effect of random errors in the measurement and comparison of gate output energies in each time slot. Suppose that measurements of the optical energies emerging from gates 1 to 4, integrated over the duration of a burst, produce the results E_1 to E_4 respectively. The difference $E_i - E_j$ (for all $i, j = 1, \dots, 4$, provided $j \neq i$) is assumed to be a normally-distributed random variable with mean $\langle E_{ij} \rangle = \langle E_i \rangle - \langle E_j \rangle$ and variance σ_c^2 . Then, for example, the probability $P(\langle E_1 \rangle < \langle E_2 \rangle)$ is equal to $\Phi(\langle E_{21} \rangle / \sigma_c)$, where $\Phi(z)$ is given by (4). Therefore the probability

$P(\langle E_1 \rangle = \min \{E_i\}) = P(\langle E_1 \rangle < \langle E_2 \rangle)P(\langle E_1 \rangle < \langle E_3 \rangle)P(\langle E_1 \rangle < \langle E_4 \rangle)$ is equal to

$\Phi(\langle E_{21} \rangle / \sigma_c)\Phi(\langle E_{31} \rangle / \sigma_c)\Phi(\langle E_{41} \rangle / \sigma_c)$. More generally,

$$P(\langle E_j \rangle = \min \{E_i\}) = \prod_{i \neq j} \Phi(\langle E_{ij} \rangle / \sigma_c) , \quad (16)$$

and this, according to the strategy (15), is the probability that gate $J = 1 + (j+1) \bmod 4$ will be selected. If gate J is selected, then the bit-error probability is $\langle e | \psi_J \rangle$, as given by (9) with ψ_J given by (14). It follows that the overall bit-error probability is

$$\langle e | \theta \rangle = \sum_{j=1}^4 \langle e | \psi_j \rangle \prod_{i \neq j} \Phi(\langle E_{ij} \rangle / \sigma_c) . \quad (17)$$

The curves shown in Fig. 12 with $\sigma_c > 0$ illustrate the effect of random errors in the gate-selection mechanism for the quad-gate regenerator. A measurement standard deviation $\sigma_c = 0.1\%$ is found to have a negligible effect on the bit-error probability $\langle e \rangle$ (unlike in the case of the dual-gate regenerator where $\sigma_c = 0.1\%$ was sufficient to increase $\langle e \rangle$ by more than 9 orders of magnitude, as shown in Fig. 9). Indeed, in this example only when the standard deviation σ_c becomes as large as $\sim 5\%$ is there a noticeable increase in bit-error probability for the quad-gate regenerator, demonstrating a robust tolerance of random errors in the gate-selection mechanism. This is because the most likely outcome of an incorrect selection is to choose the gate closest to the optimum one, rather one of the other outputs which may contain large numbers of errors.

Figure 13(a) shows an example of the worst-case bit-error probability, $\max\langle e | \theta \rangle$, for any phase angle θ in $[0, 2\pi]$. It is found that with $W = 0.75T$, $\sigma \leq 0.032T$ and $\sigma_c \leq 5\%$, $\max\langle e | \theta \rangle$ is expected to be less than 10^{-12} .

D. Multi-Gate Regenerators

Asynchronous regenerators with larger numbers of gates are expected to exhibit greater jitter tolerance, but at the price of greater complexity. In the limit of a very large number of gates, each with $W \rightarrow T$, and ultrashort clock pulses that can be

represented as impulses, the jitter tolerance is expected to be twice that of the quad-gate regenerator (i.e. a maximum bit-error probability of 10^{-12} is predicted for rms jitter $\sigma \leq 0.071T$). This is identical to the predicted optimum performance of a bit-synchronous channel-demultiplexer for OTDM systems, in which $W = T$ [24]. We consider the 50% reduction in jitter margin obtained when using a quad-gate asynchronous regenerator (compared to a multi-gate asynchronous regenerator or a synchronous regenerator) to be an acceptable price to gain both the relative simplicity of the quad-gate design and the important advantages of asynchronous regeneration.

E. Transmission Through a Sequence of Regenerators

When a burst travels across a network it may pass through many regenerators, and the performance is determined by the end-to-end bit-error probability. As the burst travels, the phase angle θ differs from one regenerator to the next in a random fashion, because the local clocks at the regenerators are independent and therefore uncorrelated. In the Appendix the bit-error probability is calculated for a data burst that has passed through a sequence of Q identical regenerators, with the assumption that the rms jitter σ in the arrival time of the data bits relative to the local clock is the same at each regenerator. Figure 13(b) is a plot of the expected bit-error probability $\langle e \rangle$ for a given burst, per regenerator passed, according to (A2). By comparing the two curves in Fig. 13 it can be verified that, for a given rms jitter σ , the total bit-error probability $Q\langle e \rangle$ after passing Q regenerators with a uniform random distribution of phase angles θ is significantly smaller than Q times the worst-case bit-error probability, $\max\langle e | \theta \rangle$, for each individual regenerator.

The result shown in Fig. 13(b) applies to the transmission of an isolated burst across the network. However this must be treated with caution when considering the transmission of a string of bursts from the same source. This is because, although the phase θ of a data burst relative to the local clock at a regenerator is a continuous uniform random variable, nevertheless the phases of successive bursts arriving at the regenerator in a string from the same source may be highly correlated. This is because in some cases the main physical fluctuations that cause θ to vary (such as temperature and mechanical changes acting on an optical fibre cable) may occur on a slower time scale than the duration of the transmitted string of bursts. Therefore it

may happen that the string, after transmission across the network, contains a significant number of errors, even though the bit-error probability according to (A2) may be very low. Therefore when designing a network containing many regenerators, it may be appropriate to take a more conservative estimate of performance based on the worst-case bit-error probability, $\max\{e|\theta\}$, at each regenerator.

F. Burst Time-Position Slippage

A side-effect of the asynchronous regeneration we have described is that it introduces a small non-deterministic variation in the time-of-flight of a burst through the regenerator. This delay variation can be written as $d = \tau + \delta$, where τ is the displacement of the regenerated burst resulting from the phase difference θ between the incoming burst and the local clock, and δ is the path-delay variation that can occur if the optical paths through the regenerator via the different gates are not exactly equalised. Whereas the delay variation δ may be minimised by careful design and construction, the displacement τ is an inherent feature of this type of regenerator, and is bounded on the interval $[-T/2k, +T/2k]$ where k is the number of gates. In the case of a single regenerator these small random time displacements of the regenerated bursts may be considered negligible because they are only a fraction of the bit period T . However a burst that travels across a network through a sequence of many regenerators may accumulate a significant net time displacement. A system error may occur if the accumulated displacement exceeds some specified limit of L bit periods (which may be an appropriate fraction of the guard band shown in Fig. 1).

This accumulated displacement in time experienced by a burst is similar to the problem of a random walk, or the Brownian motion of a particle, in one dimension. Assuming the path-delay variation δ can be neglected, at the i th regenerator the regenerated burst undergoes a non-deterministic time shift τ_i relative to the average delay. To a good approximation, τ_i can be considered to follow a continuous uniform random distribution, $\tau_i \sim U[-T/2k, +T/2k]$, with probability density function $U(\tau_i) = k/T$. After a burst has passed through Q identical regenerators (with uncorrelated phases), its accumulated displacement from the average delay value is $\sum \tau_i$, which has an expected value of zero and variance $QT^2/12k^2$. The probability

that the magnitude of the accumulated displacement exceeds the specified limit of LT is therefore given by

$$P(|\sum \tau_i| > LT) = 2P\left(Z > \frac{LT}{\sqrt{Q/12}}\right), \quad (18)$$

where $Z \sim N(0,1)$ is the standard normal random variable, and it is assumed that Q is large (i.e. $Q > \sim 50$). For small values of Q we can use the property that

$P(|\sum \tau_i| > LT)$ is exactly zero for $Q < 8L$ (since a directed (non-random) walk consisting of $8L$ steps, each of length not more than $T/8$, cannot reach a distance greater than LT).

Figure 14 shows values for the probability that a burst suffers 'slippage' (i.e. a time displacement greater than the specified limit LT), according to (18). The slippage probability is plotted versus Q , the number of regenerators passed, for various values of LT . For networks in which $Q \leq 100$, the burst-slippage probability is expected to be less than 10^{-9} if LT is set to be 5 bit periods, whereas in larger networks with $Q \approx 1000$, LT should be set to at least 14 bit periods. For a burst size of $n \geq 1000$ bits, these values represent only a small overhead per time slot.

IV. AN EXAMPLE NETWORK USING DIGITAL OPTICAL REGENERATION

According to our definition given earlier, an asynchronous digital photonic network is one in which the network nodes operate without global synchronisation at the bit level. A simple example of how digital optical regeneration may be used in such a network is shown in Fig. 15. It will be noticed that digital regenerators of both the conventional bit-synchronous (SR) type (i.e. using clock recovery on a continuous data stream) and the bit-asynchronous (AR) type (i.e. of the type described here, not using clock recovery) can be used to advantage. In the example shown, three sources of bursts are depicted (A, B and C), each of which have independent clocks. The sources may operate at the same or different wavelengths. The outgoing link from each source carries only bursts that are in bit-synchronism with the clock in that source and therefore with each other. This means that if a regenerator is required in such a link, it may be a SR in which the local pulse source is forced into synchronism with the incoming data (using clock recovery). However, in general the inputs to the

routing and processing nodes will be bit-asynchronous. Thus, for example in Fig. 15, bursts that arrive at node 1 having originated from sources A and B are not phase coherent because the clocks in sources A and B are uncorrelated. Each input to the network node is regenerated using an AR, and each AR within the node shares the same local clock pulse source. Depending on the routing function within the node, each output link may carry bursts that have originated from more than one source (for example, in Fig. 15, an output link from node 1 contains bursts that have originated from A and B). However, with careful design (i.e. by equalising the optical paths inside the regenerator), the asynchronous regenerators AR in the node can ensure that all the bursts transmitted on an output link, regardless of origin, are in bit-synchronism with each other and with the local clock in the node. Therefore if a regenerator is required further down this link, it may again be of SR type. It might also be advantageous in some cases to use a SR at an input to a node immediately before the AR. This would serve to suppress jitter and other signal impairments, and thus relax the design constraints for the AR. By selecting the wavelength of the local clock sources appropriately, perhaps in combination with ultrafast wavelength conversion, optimum use can be made of the available spectrum on the fibre infrastructure used to connect network nodes.

V. CONCLUSIONS

This paper has outlined the concept of the asynchronous digital photonic network. The network is based on digital optical transmission and processing, which provides a modular and cascadeable approach to network element design and the prospect of almost-infinite extendibility of the physical network. Optical transmission throughout the network is in burst mode, using a standard optical signal format. The main defining feature of the asynchronous digital photonic network is that the network nodes (in which the optical processing and routing take place) do not share global synchronisation at the bit level. Instead, the nodes each operate with independent bit-level clocks.

An essential sub-system in the network nodes is the asynchronous digital optical regenerator. This regenerates incoming bursts, but unlike a traditional synchronous regenerator, the local pulse source is not forced into synchronism with the incoming

bits. Instead, the asynchronous regenerator introduced here works in a reverse fashion, by forcing the regenerated bursts into synchronism with an independent local clock. Some significant technical advantages of this type of regeneration have been mentioned: it allows a free-running mode-locked laser or a laser with external pulse compression to be used as the local clock at a node, which provides the highest-quality optical pulses for processing and onward transmission; and this clock can be used as a continuous supply of synchronised pulses for subsequent bit-level digital optical processing within the node. The most important feature of the asynchronous regenerator is that, as well as providing 3R digital regeneration, it re-synchronises all the incoming traffic at a node to the local bit-level clock. Various design and performance aspects of this new type of regenerator have been described. A simple example of a network implementation using these principles has also been given.

By considering this and further examples we reach an interesting and important observation. With suitable design, every link in the asynchronous network may be arranged to carry exclusively bit-synchronous traffic (i.e. traffic which, regardless of origin, is synchronised to the same local bit-level clock). Moreover all the traffic on a link shares a defined digital optical format (i.e. wavelength, pulse duration, intensity, etc). This greatly simplifies the design of a network on 'digital' modular principles (for example, by removing the need for burst-by-burst power level equalisation prior to processing within the nodes), and also allows the prospect of ultra-broadband digital regeneration using ultrafast photonic devices. But most important, by regenerating the bursts at each node in bit-asynchronous fashion, this eliminates the need to maintain bit-level synchronism between the different links and routing nodes throughout the network.

As stated in the Introduction, the aim of developing the concept of the asynchronous digital photonic network is to combine the advantages, and overcome the drawbacks, of the various existing approaches to photonic networking. The asynchronous digital photonic network uses high-speed digital processing in the optical domain to achieve almost infinite scalability, but avoids the need for global bit-level synchronisation which would otherwise severely constrain the network architecture.

Finally, we comment that experimental proof-of-principle demonstrations of the concepts described here are currently in progress in our laboratories.

APPENDIX:

BIT-ERROR PROBABILITY FOR TRANSMISSION THROUGH A SEQUENCE OF ASYNCHRONOUS REGENERATORS

We consider a burst that passes through a sequence of asynchronous regenerators, each of which has an independent clock. Therefore, at each regenerator, θ is a continuous uniform random variable, $\theta \sim U[0, 2\pi]$. We assume here that the regenerators each have k gates and that the rms jitter σ in the arrival time of the data bits relative to the local clock is the same at each regenerator. It follows that the bit-error probability after passing through Q regenerators is given by $\mathcal{Q}(e) = \mathcal{Q}(\langle e|\theta \rangle)$, where $\langle e|\theta \rangle$ is the conditional bit-error probability for a single regenerator. By inspection, $\langle e|\theta \rangle$ is a periodic, symmetrical function of θ that exercises k full periods on $[0, 2\pi]$ (for example, see Fig. 7 for $k = 2$ and Fig. 11 for $k = 4$). A minimum of $\langle e|\theta \rangle$ occurs when the first gate window is positioned symmetrically over the clock pulse, i.e. when $\theta/2\pi = W/2T$. The next maximum of $\langle e|\theta \rangle$ occurs when $\theta/2\pi = W/2T + 1/2k$. The interval between these two points (i.e. $W/2T < \theta/2\pi < W/2T + 1/2k$) represents one-half period of $\langle e|\theta \rangle$. It follows that the bit-error probability $\mathcal{Q}(e) = \mathcal{Q}(\langle e|\theta \rangle)$ with $\theta \sim U[0, 2\pi]$ is identical to the expected value $\mathcal{Q}(\langle e|\theta \rangle)$ with $\theta/2\pi \sim U[W/2T, W/2T + 1/2k]$. Within this latter interval, $\langle e|\theta \rangle$ is a continuous and differentiable function given by (9), and θ has the expected value $\langle \theta \rangle = \pi(W/T + 1/2k)$ and $\text{var } \theta = \pi^2/24k^2$. Now let $Y = \ln\langle e|\langle \theta \rangle + \Delta \rangle$ and perform a Taylor expansion:

$$Y = \ln\langle e|\langle \theta \rangle \rangle + \Delta \frac{d}{d\theta} \ln\langle e|\theta \rangle \bigg|_{\langle \theta \rangle} + \frac{\Delta^2}{2!} \frac{d^2}{d\theta^2} \ln\langle e|\theta \rangle \bigg|_{\langle \theta \rangle} + \frac{\Delta^3}{3!} \frac{d^3}{d\theta^3} \ln\langle e|\theta \rangle \bigg|_{\langle \theta \rangle} + \dots \quad (A1)$$

Taking expectations to find $\langle Y \rangle = \ln\langle e \rangle$, we obtain

$$\mathcal{Q}(e) \approx \mathcal{Q} \exp \left[\ln\langle e|\langle \theta \rangle \rangle + \frac{\text{var } \theta}{2} \frac{d^2}{d\theta^2} \ln\langle e|\theta \rangle \bigg|_{\langle \theta \rangle} \right], \quad (A2)$$

since $\langle \Delta \rangle = 0$, $\langle \Delta^2 \rangle = \text{var } \theta$ and $\langle \Delta^3 \rangle = 0$. To evaluate (A1) we use the following relations derived from (9):

$$\langle e|\theta \rangle \approx \frac{1}{2} \left[1 - \Phi\left(\frac{\theta/2\pi}{\sigma/T}\right) + \Phi\left(\frac{\theta/2\pi - W/T}{\sigma/T}\right) \right], \quad (\text{A3})$$

$$\frac{d}{d\theta} \langle e|\theta \rangle \approx \frac{\pi}{\sigma/T} \left[\phi\left(\frac{\theta/2\pi - W/T}{\sigma/T}\right) - \phi\left(\frac{\theta/2\pi}{\sigma/T}\right) \right], \quad (\text{A4})$$

$$\frac{d^2}{d\theta^2} \langle e|\theta \rangle \approx \frac{2\pi^2}{(\sigma/T)^3} \left[\frac{\theta}{2\pi} \phi\left(\frac{\theta/2\pi}{\sigma/T}\right) - \left(\frac{\theta}{2\pi} - \frac{W}{T}\right) \phi\left(\frac{\theta/2\pi - W/T}{\sigma/T}\right) \right], \quad (\text{A5})$$

$$\frac{d^2}{d\theta^2} \ln \langle e|\theta \rangle = \frac{1}{\langle e|\theta \rangle} \frac{d^2}{d\theta^2} \langle e|\theta \rangle - \frac{1}{\langle e|\theta \rangle^2} \left(\frac{d}{d\theta} \langle e|\theta \rangle \right)^2, \quad (\text{A6})$$

where $\Phi(z)$ and $\phi(z)$ are given by (7) and (8). Results for the case with $k = 4$ are shown in Fig. 13(b).

ACKNOWLEDGEMENTS

The authors are grateful to Sarah Cotter and Julian Lucek for useful discussions.

REFERENCES

- [1] A. D. Ellis, T. Widdowson, X. Shan and D. G. Moodie, "Three node, 40 Gbit/s OTDM network experiment using electro-optic switches," *Electron. Lett.*, vol. 30, pp. 1333–1334, 1994.
- [2] A. D. Ellis and T. Widdowson, "690 node global OTDM network demonstration," *Electron. Lett.*, vol. 31, pp. 1171–1172, 1995.
- [3] T. Matsumoto, M. Jinno, K. Kimura and K. Noguchi, "Studies on optical digital cross-connect systems for very high speed optical communications networks," *Proc. ICC'94*, LOCATION, pp. 1060–XXXX, 1994.
- [4] S.-W. Seo, K. Bergman and P. R. Prucnal, "Transparent optical networks with time-division multiplexing," *IEEE J. Selected Areas. Comm.*, vol. 14, pp. 1039–1051, 1996.

- [5] D. Cotter, J. K. Lucek and D. D. Marcenac, "Ultra-high bit rate networking: From the transcontinental backbone to the desktop", *IEEE Commun. Mag.*, vol. 35, pp. 90-95, 1997.
- [6] A. D. Ellis, T. Widdowson, I. D. Phillips and W. A. Pender, "High speed OTDM networks employing electro-optic modulation", *IEICE Trans. Electron.*, vol. E-81C, no. 8 (in press), 1998.
- [7] R. S. Vodhanel *et al.*, "Performance of an 8 wavelength 8 node WDM ring experiment with 80 Gbit/s capacity," *Proc. OFC'96*, 1996, paper PD28.
- [8] V. Mizrahi *et al.*, "Future WDM systems," *Proc. ECOC'97*, Edinburgh, 1997, paper MO4A1.
- [9] M. Jinno and M. Abe, "All optical regenerator based on nonlinear fibre Sagnac interferometer," *Electron. Lett.*, vol. 28, pp. 1350-XXXX, 1992.
- [10] J. K. Lucek and K. Smith, "All optical signal regenerator," *Opt. Lett.*, vol. 18, pp. 1226-XXXX, 1993.
- [11] W. A. Pender, T. Widdowson and A. D. Ellis, "Error free operation of a 40 Gbit/s all optical regenerator," *Electron. Lett.*, vol. 32, pp. 567-569, 1996.
- [12] D. Cotter *et al.*, "Self-routing of 100 Gbit/s packets using 6 bit 'keyword' address recognition," *Electron. Lett.*, vol. 31, pp. 1475-1476, 1995; publisher's addendum vol. 31, pp. 2201-2202, 1995.
- [13] M. Steenstrup (ed.), *Routing in Communications Networks*. Englewood Cliffs, NJ: Prentice Hall, 1995, and references therein.
- [14] F. Borgonovo, *ibid.*, chapter 9.
- [15] P. R. Prucnal, "Optically processed self-routing, synchronisation and contention resolution for 1D and 2D photonic switch architectures," *IEEE J. Quantum Electron.*, vol. 29, pp. 600-612, 1993.
- [16] M. Sexton and A. Reid, *Transmission Networking: SONET and the Synchronous Digital Hierarchy*. Boston, MA; London: Artech House, 1992.

- [17] T. Widdowson, D. J. Malyon, X. Shan, and P. J. Watkinson, "Soliton propagation without transmission control using a phase locked erbium fibre ring laser," *Electron. Lett.*, vol. 30, pp. 661–663, 1994.
- [18] T. Morioka, S. Kawanishi, K. Mori and M. Sarawutari, "Transform-limited, femtosecond WDM pulse generation by spectral filtering of gigahertz continuum," *Electron. Lett.*, vol. 30, pp. 1166–1167, 1994.
- [19] M. Nakazawa, E. Yoshida, H. Kubota and Y. Kimura, "Generation of a 170 fs, 10 Ghz transform-limited pulse train at 1.55 μ m using a dispersion-decreasing, erbium-doped active soliton compressor," *Electron. Lett.*, vol. 30, pp. 2038–2040, 1994.
- [20] D. J. Richardson *et al.*, "Demonstration of 205 km transmission of 35 GHz, 5 ps pulses generated from a diode driven, low jitter, beat signal to soliton train conversion source," *Electron. Lett.*, vol. 31, pp. 470–472, 1995.
- [21] K. J. Blow, R. J. Manning and A. J. Poustie, "All-optical bit serial processing," in A. Sasaki (ed.), *Optical Interconnect and Computing*. Oxford: Oxford University Press, 1999 (in preparation).
- [22] N. J. Doran and D. Wood, "Nonlinear-optical loop mirror", *Opt. Lett.*, vol. 13, pp. 56–58, 1988.
- [23] N. A. Whitaker *et al*, *Opt. Lett.*, vol. 16, pp. 1838–1840, 1991.
- [24] M. Jinno, "Effects of crosstalk and timing jitter on all-optical time-division demultiplexing using a nonlinear fibre sagnac interferometer switch," *IEEE J. Quantum Electron.*, vol. 30, pp. 2842–2853, 1994.
- [25] K. Kang, I. Glesk and P. R. Prucnal, "Ultrafast optical time demultiplexers using semiconductor optical amplifiers," *Inter. J. High Speed Electron. & Systems*, vol. 7, pp. 125–151, 1996.
- [26] R. J. Manning, A. D. Ellis, A. J. Poustie and K. J. Blow, "Semiconductor laser amplifiers for ultrafast all-optical signal processing", *J. Opt. Soc. Am. B*, vol. 14, pp. 3204–3216, 1997.

- [27] I. D. Phillips *et al.*, "Drop and insert multiplexing with simultaneous clock recovery using an electroabsorption modulator," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 291–293, 1998.
- [28] A. Bononi, F. Forghieri and P. R. Prucnal, "Soliton ultrafast all-optical networks," *IEE Proc.-J*, vol. 140, 285–290, 1993.

Data bit			Bit-error probability
$j-1$	j	$j+1$	
0	0	0	0
0	0	1	$q/8$
0	1	0	$(1-p)/8$
0	1	1	$[1-p-q+(2-\alpha)pq]/8$
1	0	0	$r/8$
1	0	1	$[q+r-(2-\alpha)qr]/8$
1	1	0	$[1-p-r+(2-\alpha)pr]/8$
1	1	1	$[(1-p)(1-q)(1-r) + (1-\alpha)(pq+pr+qr-3pqr)]/8$

Table 1: Bit-error probability in the j th bit position at the output from a gate, for various input bit patterns (assuming $0 < \psi < W/T$).

Data bit			Probability of a 'mark'
$j-1$	j	$j+1$	
0	0	0	0
0	0	1	$q/8$
0	1	0	$p/8$
0	1	1	$[p + q - (2 - \alpha)pq]/8$
1	0	0	$r/8$
1	0	1	$[q + r - (2 - \alpha)qr]/8$
1	1	0	$[p + r - (2 - \alpha)pr]/8$
1	1	1	$[p + q + r - (2 - \alpha)(pq + pr + qr) + 3(1 - \alpha)pqr]/8$

Table 2: Probability of a 'mark' in the position of the j th clock pulse at the output from a gate, for various input bit patterns (assuming $0 < \psi < W/T$).

FIGURE CAPTIONS

Figure 1: Time slots for burst-mode transmission.

Figure 2: Digital optical regenerator for data represented by return-to-zero pulses.

Figure 3: Time slippage Δt of the local clock over the full duration of a burst, for various values of the nominal bit rate $1/T$, assuming $|\Delta f| = 50$ kHz.

Figure 4: Dual-gate asynchronous burst regenerator. The circuit consisting of detectors D1, D2 and comparator C measures and compares the energies of the burst at the output of G1 and G2, in each time slot.

Figure 5: Timing diagram for the dual-gate regenerator.

Figure 6: Probability density function of the arrival time of the i th data bit at the gates, in the time frame of the local bit-level clock.

Figure 7: Bit-error probability in the output from individual gates in the dual-gate regenerator, for the case $\sigma = 0.018T$ and $W = 0.75T$.

Figure 8: Bit-error probability for the dual-gate regenerator with ideal gate selection, for two values of rms jitter σ , and $W = 0.75T$.

Figure 9: Bit-error probability for the dual-gate regenerator when the gate-selection mechanism is subject to random errors, with $\sigma = 0.018T$, $W = 0.75T$, and $\sigma_c = 0.1\%$.

Figure 10: Quad-gate asynchronous burst regenerator.

Figure 11: Bit-error probability in the output from individual gates in the quad-gate regenerator, for the case $\sigma = 0.036T$ and $W = 0.75T$.

Figure 12: Bit-error probability for the quad-gate regenerator with $\sigma = 0.036T$ and $W = 0.75T$, in the case of ideal ($\sigma_c = 0$) and imperfect ($\sigma_c > 0$) gate selection.

Figure 13: a) Worst-case bit-error probability for a single quad-gate regenerator, and b) bit-error probability per regenerator passed for a data burst that has passed Q uncorrelated quad-gate regenerators. For both curves, $W = 0.75T$.

Figure 14: Probability that a burst will suffer 'slippage' (displacement in time by more than an allowed margin LT , where T is the bit period) after passing Q regenerators.

Figure 15: Examples of the use of digital regeneration in an asynchronous photonic network (AR = asynchronous regenerator, SR = synchronous regenerator). Bursts that have originated from source A, B or C are shown white, black or shaded, respectively.

Single Gate Regenerator

In an alternative embodiment, the local clock pulse source is again continuously free-running, but requires only one gate to modulate the output of the 5 clock pulse source so as to regenerate the packet.

The principle of this alternative approach is shown in Figure 21. In each time slot, the phase detector measures the phase angle θ between the free-running local pulse source and the incoming packet. This information is used to shift by an appropriate amount the phase of the control signal that is applied to the gate. The 10 effect of the phase shifter is that when a packet data bit of value 1 causes the gate window to open, the window is located as near as possible centrally over the clock pulse (as depicted in Figure 21). The phase detector and phase shifter operate once in each time slot.

In the ideal case the phase detector and phase shifter operate without error and 15 the gate window is located exactly centrally over the clock pulse. In that case the main source of bit errors in the regenerated packet is jitter in the arrival time of the packet data bits. The analysis of bit errors arising from timing jitter in this regenerator is then similar to the analysis of the effects of timing jitter in an OTDM demultiplexer given by Jinno (IEEE Journal of Quantum Electronics, vol 30, no.12, 20 pp. 2842-2853, 1994). In particular, the bit-error probability as a function of the ratio of the gate widow width W to the bit period T , for various values of the rms jitter σ in the arrival time of the packet data bits, is as depicted in Figure 5 of Jinno (1994). Some results of that analysis are that the minimum bit-error probability is obtained when W is equal to T , and also σ must be less than 25 $0.071T$ to ensure that the bit-error probability is less than 10^{-12} .

In a more realistic case, noise and other imperfections in the phase detector and phase shifter cause systematic and random errors in the position of the gate window relative to the clock pulse. We neglect systematic errors and assume here that, as a result of the random errors, the position of the gate window is a 30 normally distributed random variable with standard deviation σ_w . In that case, the bit-error analysis is again similar to that of Jinno (1994), and the results shown in Figure 5 of Jinno can be used, except with the parameter σ replaced by $\sqrt{\sigma^2 + \sigma_w^2}$. Again it is found that the minimum bit-error probability is obtained

when W is equal to T , and $\sqrt{\sigma^2 + \sigma_{\text{lw}}^2}$ must be less than $0.071T$ to ensure that the bit-error probability is less than 10^{-12} .

Figure 22 shows a possible embodiment of this alternative version of the bit-asynchronous regenerator, and it is assumed that the bit rate is 100 Gbit/s. The 5 incoming packet (at wavelength λ_{in}) is first passed through an optical input stage with slowly-responding automatic level control (such as an erbium-doped fibre amplifier), and is then split into two paths, one leading to the phase detector and the other to the phase shifter. The local clock pulse source is a mode-locked ring fibre laser producing ~2 ps pulses at the wavelength λ_c . The output of this clock 10 source is also split into two paths, one leading to the phase detector and the other to the optical gate. At the input to the phase detector, the clock pulses are broadened to ~10 ps, for example by passing them through the optical bandpass filter F1 (as shown in Figure 22) or by dispersion in fibre or in a chirped grating. The phase detector could be based on four-wave mixing in a semiconductor optical 15 amplifier SOA FWM (as described for example by O Kanatani, S Kawanashi and M Sarawutari in Electronics Letters, vol.30, no.10, p.807, 1994). The output from the SOA FWM is isolated using the optical bandpass filter F2, and then detected. The electronic processing stage (which could include a low noise, high linearity sample-and-hold gate triggered by the global packet-level clock) measures the 20 photodetector output voltage immediately after the arrival of the packet in each time slot, the measured voltage being given approximately by $A + B\cos\theta$, where A and B are constants and θ is the phase difference between the clock and the incoming packet bits. In the example implementation of the phase shifter, this signal is used to control the wavelength λ_{cw} of the continuous-wave distributed- 25 feedback laser DFB in the phase shifter section. The output of the DFB laser is connected to the input of a date-driven optical switch denoted UNI1. This device could be the ultrafast nonlinear interferometer switch described by Hall and Rauschenbach (paper PD5, Proceedings of Conference on Optical Fiber Communication OFC'98, published by the Optical Society of America, February 30 1998), which has been shown to operate at a speed of 100 Gbit/s, although any data driven optical switch capable of producing approximately square switching windows with negligible variation in throughput delay would be suitable. The control signal to UNI1 is the input data packet. The action of UNI1 is therefore to

shift the wavelength of the incoming packet from λ_{in} to the controlled wavelength λ_{cw} . This packet with shifted wavelength is then isolated using optical filter F3, and passed through a dispersive optical delay line which imparts a time delay which depends on wavelength. This dispersive delay line could be a fibre grating 5 device or, as shown in Figure 22, a length of optical fibre, such as a length of optical fibre of the type manufactured primarily for use in dispersion compensation. The required minimum amount of phase shift of the control pulses is ± 5 ps (i.e. $\pm 0.5T$, where T is the bit period). Therefore, if for example the length of dispersion compensating fibre is 20m and the fibre dispersion is 100 ps/nm/km, 10 the required shift in the wavelength λ_{cw} of the DFB laser is ± 2.5 nm, and this must be accomplished within the duration of the guard band between packets (say ~ 5 ns). Alternatively, instead of controlling the wavelength of a single DFB laser, as shown in Figure 22, the control loop may contain means to select one of a number of lasers each having a different fixed wavelength. Because the optical 15 fibre used as a dispersive delay line is a long device subject to changes in path length due to environmental factors (temperature, strain, etc), it is convenient as in Figure 22 to pass the local clock pulses over the same fibre. This ensures that there is a negligibly small relative change in delay for the clock and control pulses. Alternatively, the error signal could be used to select one of a discrete number of 20 optical paths.

A further alternative phase shifter comprises means to select one from a number of optical delay lines each having a different fixed delay. The delay lines could consist of a silicon wafer on which is fabricated a silica-on-silicon planar lightwave circuit. This circuit may be integrated in hybrid fashion with an array of 25 discrete or integrated semiconductor optical switching devices, such as semiconductor optical amplifiers or electroabsorption modulators. In this case the electronic processing stage selects the appropriate delay line by switching on or off the appropriate semiconductor optical switching devices.

The resultant optical data bits, suitably phase shifted, are then used as 30 the control pulses in the optical gate. The gate, denoted UNI2 in Figure 22, which is controlled by the phase-shifted packet data bits, is used to modulate the locally generated clock pulses so as to produce a regenerated packet, synchronous with the local clock. Again the device UNI2 could be the ultrafast nonlinear

interferometer switch described by Hall and Rauschenbach. In the case that the regenerator is receiving inputs from a multiple number of sources, if the embodiment shown in Figure 22 is employed, the same length of dispersion-compensating fibre should be used for all inputs to provide appropriate phase 5 shifts, so that all the regenerated packets are bit-synchronous.

There is described above with reference to Figure 21, an alternative approach to the bit-asynchronous packet regenerator, in which the local clock pulse source is again continuously free-running, but which requires only one gate to modulate the output of the clock pulse source so as to regenerate the packet. A 10 measurement of the phase θ was used to control a phase shifter acting on the control signal to the optical gate. This alternative approach is, in effect, a feed-forward, open loop control system. There is also described, in the immediately preceding section, the use of a regenerator in a network in which incoming packets at the regenerator are bit-synchronous (i.e. they all share the same bit- 15 level clock). Therefore in this case the phase θ is only very slowly varying from one packet to the next, and this allows another arrangement for the control of the phase shifter, as will be described with reference to Figures 24 and 25.

Figure 24 shows the principle of this alternative approach. In this case the phase detector measures the phase angle θ between the free-running local 20 pulse source and the control signal applied to the optical gate (i.e. the packet data bits after the phase shifter). The measured phase angle is then used to control the phase shifter. The control signal to the phase shifter may be either an analogue or digital signal (preferably a digital signal at the packet level), and this control signal would be gated at the rate of once per packet time slot. In contrast with the 25 arrangement shown in Figure 21, this is a feed-back, closed loop control system. This has the advantage that the system is free of systematic errors and drift, even if the phase detector is nonlinear. Furthermore, drifts in the phase shifter are automatically compensated for by the closed loop (because the phase shifter is inside the feedback path). Because, in a practical system, the feedback delay may 30 be greater than the duration of a packet time slot, this closed-loop control system will not be sufficiently fast acting to track substantial variations in phase θ from packet to packet. However this is not an important limitation in certain important network applications, as described above.

Figure 25 shows an example embodiment of this alternative form of the bit-asynchronous packet regenerator. The various designated components are as described previously for Figure 22.

The description above includes a discussion of the allowable amount of 5 frequency difference between the clock at the packet source and the clock in the asynchronous regenerator. The present embodiment of the asynchronous regenerator using a feed-back, closed-loop arrangement may impose a further restriction on the amount of frequency difference that can be tolerated. It is necessary that the frequency difference between the bit rate of the incoming 10 packet and the full-rate optical clock source in the regenerator is significantly smaller than the effective bandwidth of the control loop (including the electronic bandwidth, the feedback delay and the speed of response of the phase shifter). Following normal engineering practice, the frequency offset should be at least an order of magnitude smaller than the effective bandwidth of the control loop. For 15 example, if the determining factor for the control-loop bandwidth is an electronic bandwidth of 10 kHz, then the magnitude of the frequency offset $f_R - f_S$ should be no greater than 100 Hz (where f_S and f_R are the frequencies of the microwave oscillators depicted in Figure 4, and it is assumed that $M_S = M_R = 10$).

In an alternative embodiment, the phase shifter shown in Figure 25, 20 consisting of DFB laser, switching device UNI1, optical filter F3 and dispersion-compensating fibre, could be replaced by a variable optical delay line. Since, in the network scenario discussed above, and illustrated by Figure 23, the packets arriving at the regenerator may be in bit-synchronism. Therefore the bit-asynchronous regenerator merely needs to track the relatively slow variations in 25 the phase difference between the incoming packets and the local clock, rather than abrupt packet-to-packet phase variations. The control loop may therefore be relatively slow acting (much slower than on a packet-by-packet basis). However, the control loop bandwidth should not be so low as to restrict the amount of frequency offset that can be accommodated. Suppose, for example, that the delay 30 line is a variable motor-controlled device, such as a motor-driven fibre stretcher, which is capable of changing the value of the optical delay at a maximum rate of 100 ps per second. This corresponds to 10 bit periods per second at 100 Gbit/s, for example, and so the maximum allowable frequency offset would be an order of

magnitude less than that (to ensure the effective control loop bandwidth is at least 10 times faster than the fastest variations in the signal to be controlled), i.e. ~1 Hz, which is a severe restriction. Therefore a motor-controlled phase shifter may not have sufficient speed of response for this application. Another type of variable

5 optical delay line is a fibre stretcher consisting of a length of fibre coiled tightly around a piezo-electric drum. This type of stretcher is capable of 100 μm length change at 20 kHz, or approximately 1 ps delay change in 50 ms. In order to achieve a range of ± 5 ps (i.e. ± 0.5 bit periods at 100 Gbit/s), a cascade of piezo-electric drum stretcher units (10-20 units) would allow a frequency offset of a few

10 kHz between the local and distant clocks.

CLAIMS

1. A method of operating a node in an optical communications network including
 - a) receiving at the node an optical packet; and
 - 5 b) generating from the said optical packet received at the said node a regenerated optical packet having a phase determined by a local clock source and independent of the phase of the said packet received at the node.
2. A method of operating an optical regenerator comprising:
 - 10 a) receiving an optical packet at an input of the regenerator; and
 - b) generating from the said optical packet a regenerated optical packet having a phase determined by a local clock source and independent of the phase of the said packet received at the node.
- 15 3. A method according to claim 1 or 2, in which the step of generating a regenerated optical packet includes gating, using the received optical packet, an optical clock signal from the local clock source.
4. A method according to claim 3, including:
 - 20 i) measuring the phase of the said optical packet;
 - ii) depending on the result of step (i), modifying the phase of the optical packet; and
 - iii) subsequently applying data signals from the optical packet as a control signal to gate means arranged to gate the said optical clock signal.
- 25 5. A method according to claim 4, in which in step (i) comprises measuring the phase difference between the incoming optical packet and the local clock source.
6. A method according to claim 4 or 5, in which the gate window is equal to the
30 bit period.
7. A method according to any one of claims 4 to 6, in which the phase difference between the optical packet and the local optical clock is detected by means of a nonlinear interaction between the clock signal and the optical packet,

8. A method according to claim 7, in which the nonlinear interaction occurs in an optical fibre device.
- 5 9. A method according to claim 7, in which the nonlinear interaction occurs in a semiconductor device.
- 10 10. A method according to claim 9, in which the nonlinear interaction is a process of four-wave mixing.
11. A method according to any one of claims 4 to 10, in which the phase of the optical packet is modified by passing through a wavelength converter and a dispersive optical delay line.
- 15 12. A method according to claim 3, including:
 - passing the optical clock signal through each of a plurality of gate means;
 - applying data signals from the received optical packet as a control signal to each of the plurality of gate means with different delays of a fraction of a bit period relative to the optical clock signal input to the gate means; and
 - selecting the output of one of the plurality of gate means to provide the regenerated optical packet.
13. A method according to claim 12, in which the difference in delays is equal to T/k where T is the bit period and k is the number of optical gate means.
- 25 14. A method according to claim 12 or 13, in which the width W of the gate window is not less than T/k and not more than T , where T is the bit period and k is the number of optical gate means.
- 30 15. A method according to any one of claims 12 to 14, including making a measurement of a parameter of an optical signal output from the gate means, and selecting the output of one of the plurality of gates to provide the regenerated optical packet depending on the results of the said measurement.

16. A method according to claim 15, in which the said parameter is the energy of the optical signal.
- 5 17. A method according to claim 16, including comparing the energies of the signals from the plurality of gate means and making the said selection depending on the results of the said comparison.
- 10 18. A method according to claim 15, in which the said parameter is derived from the number of bit errors in the optical signal, and the signal with the lowest number of bit errors is selected to provide the regenerated optical packet.
- 15 19. A method according to any one of the preceding claims, including further processing the regenerated optical packet in optical processing means clocked by a signal from the local optical clock source.
20. A method of operating a communications network comprising a plurality nodes interconnected by an optical transmission medium, the method including:
 - transmitting an optical packet onto the network,
 - 20 and at a network node, receiving the said packet and generating from the said packet a regenerated optical packet having a phase determined by a local optical clock source and independent of the phase of the said packet received at the network node.
- 25 21. A method according to claim 20, including receiving at the network node optical packets from a plurality of different sources and having different respective phases.
- 30 22. A method according to claim 20 or 21, further comprising outputting the regenerated optical packet onto the optical transmission medium.
23. A regenerator for optical packets including a local optical pulse generator comprising a free-running oscillator independent in frequency and phase from the packet source.

24. A regenerator for optical packets comprising:

- means for receiving an optical packet; and
- means for generating from the said optical packet received at the said node a regenerated optical packet having a phase determined by a local clock source and independent of the phase of the said packet received at the node.

5

25. A regenerator according to claim 23 or 24, including gate means controlled by a data signal from an optical packet and connected to a local optical clock source.

10

26. A regenerator according to claim 25, including

 a plurality of gate means each arranged to receive a clock signal from a local optical source and a control signal from a packet received at the regenerator;

 means for generating different delays of the control signals relative to the

15

 clock signals at different respective gate means;

 means for selecting an output from one of the plurality of gate means.

27. A regenerator according to claim 26 including at least four gate means.

20

28. A regenerator according to claim 27, in which the ratio W/T of the gate window W to the bit period T lies substantially in the range 0.7 to 0.85.

29. A regenerator according to any one of claims 14 to 18, in which the local optical clock source is a mode-locked laser.

25

30. A regenerator according to claim 29, in which the mode-locked laser is passively mode-locked.

31. A node for connection in an optical communications network and including a

30

regenerator according to any one of claims 24 to 30

32. A node for connection in an optical communications network and arranged to operate by a method according to any one of claims 1 to 22.

33. An optical communications network including a node according to claim 31 or

32.

34. A regenerator according to claim 24 including means for measuring the phase of the optical packet, and means responsive to the said means for measuring for modifying the phase of a control signal applied to gate means, which gate means gate the optical clock signal thereby producing the regenerated optical packet.

5

35. A regenerator for optical packets according to claim 23, including an optical gate controlled by the data bits in the incoming packet, and the output of the local optical pulse generator is connected to the input of the said optical gate, and means to detect the phase difference between the incoming packet and the local optical pulse generator, and means to shift the phase of the control signal applied to the optical gate in correspondence with the said detected phase difference in such fashion as to obtain a correctly regenerated optical data packet at the output of the said optical gate.

10 15 36. A regenerator according to claim 34 or 35, in which the gate window width is equal to the bit period.

20 37. A regenerator according to any one of claims 34 to 36, in which the phase difference is detected by means of a nonlinear optical interaction between the local optical pulse generator and the incoming data packet.

25 38 A regenerator according to claim 37, in which the phase difference is detected by means of a nonlinear optical interaction between the local optical pulse generator and the incoming data packet and said nonlinear optical interaction occurs in a fibre device.

30 39 A regenerator according to claim 37, in which the phase difference is detected by means of a nonlinear optical interaction between the local optical pulse generator and the incoming data packet and said nonlinear optical interaction occurs in a semiconductor device.

40. A regenerator according to claim 39, in which the said nonlinear optical interaction is four wave mixing.

41. A regenerator according to any one of claims 34 to 40, in which the means to shift the phase of the control signal includes an optical path having a variable delay.
42. A regenerator according to any one of claims 34 to 40, in which the means to shift the phase of the control signal consists of a plurality of optical delay lines having different respective delays, and means for switching the control signal through a selected one of the plurality of optical delay lines.
43. A regenerator according to any one of claims 34 to 40, in which the means to shift the phase of the control signal consists of a wavelength converter and a dispersive optical delay line.
44. A regenerator according to claim 43, in which the said dispersive optical delay line consists of an optical fibre.
45. A regenerator according to claim 43, in which the said dispersive optical delay line consists of a fibre grating device.

15

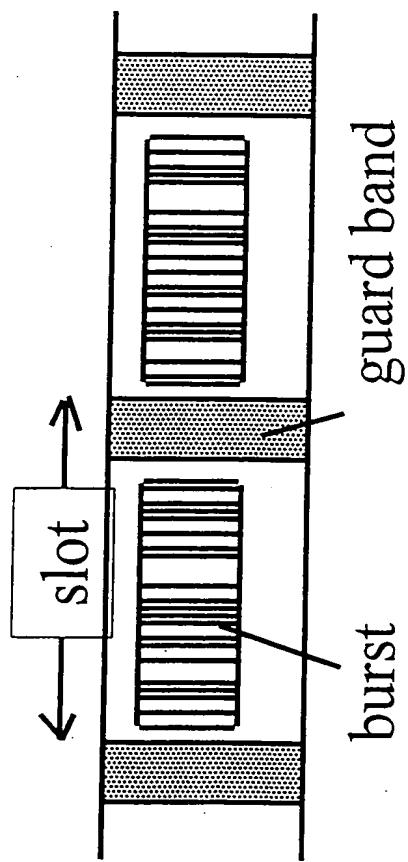
46. A regenerator according to any one of claims 34 to 45 including a feed-forward open loop control system in which the phase of the incoming packet is detected before passing through the phase shifter.
- 20 47. A regenerator according to any one of claims 34 to 45 including a feed-back closed loop control system in which the phase of the incoming packet is detected after passing through the phase shifter.
- 25 48. A regenerator according to claim 29, in which the mode-locked laser is a ring laser.
49. A node for an optical communications network including a regenerator according to any one of claims 34 to 48.
50. An optical communications network including a node according to claim 49.
- 30 51. An optical network in which bit-asynchronous regenerators are located at switching nodes.
52. An optical communications network according to claim 33, or 50 or 51 in which the links between nodes carry packets in a bit-synchronous fashion.

53. A network according to claim 52, including bit-synchronous regenerators in links between nodes.
54. A method according to any one of claims 4 to 10, in which the phase of the optical packet is modified by passing the optical packet through an optical path 5 having a variable delay.
55. A method according to any one of claims 4 to 10, in which the phase of the optical packet is modified by passing the optical packet through a selected optical delay path.
56. An optical communications network in which nodes at different locations in 10 the optical communications network process optical packets asynchronously at the bit-level.

ABSTRACT

An optical regenerator, for example at a node in an optical communications network, uses a free-running bit-asynchronous local clock source. An incoming 5 packet may be used to gate the clock source. In some implementations, a number of gates are used with different delays on the control input, and the gate giving a correctly regenerated output is selected. In alternative embodiments, a control loop is used to adjust the phase of control signals applied to a gate.

Figure 1



This Page Blank (uspto)

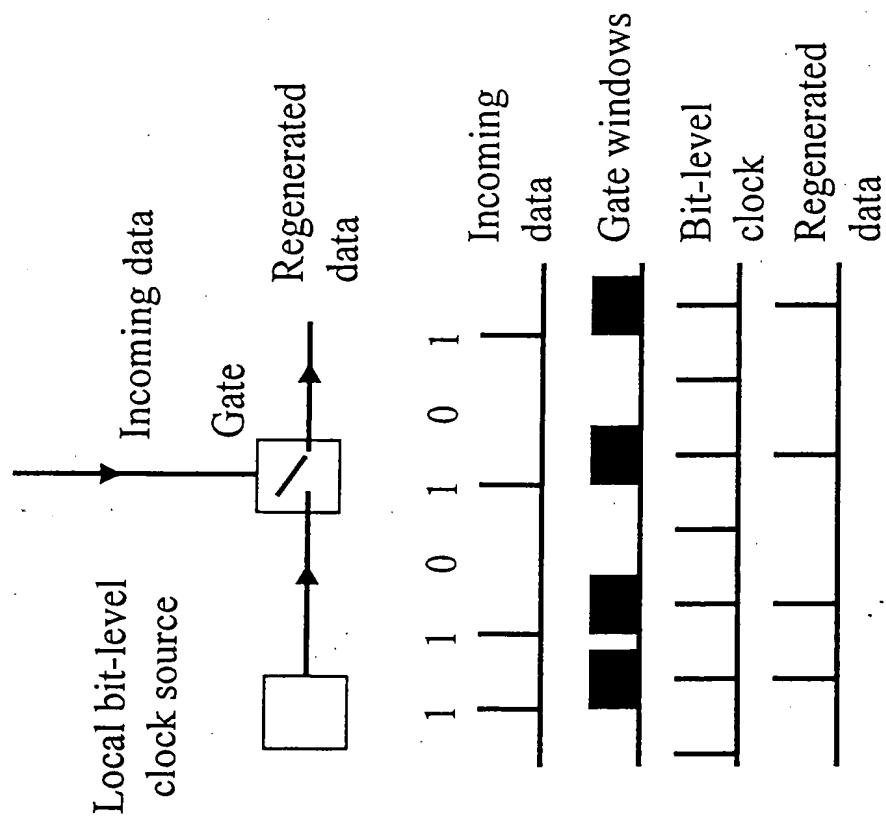
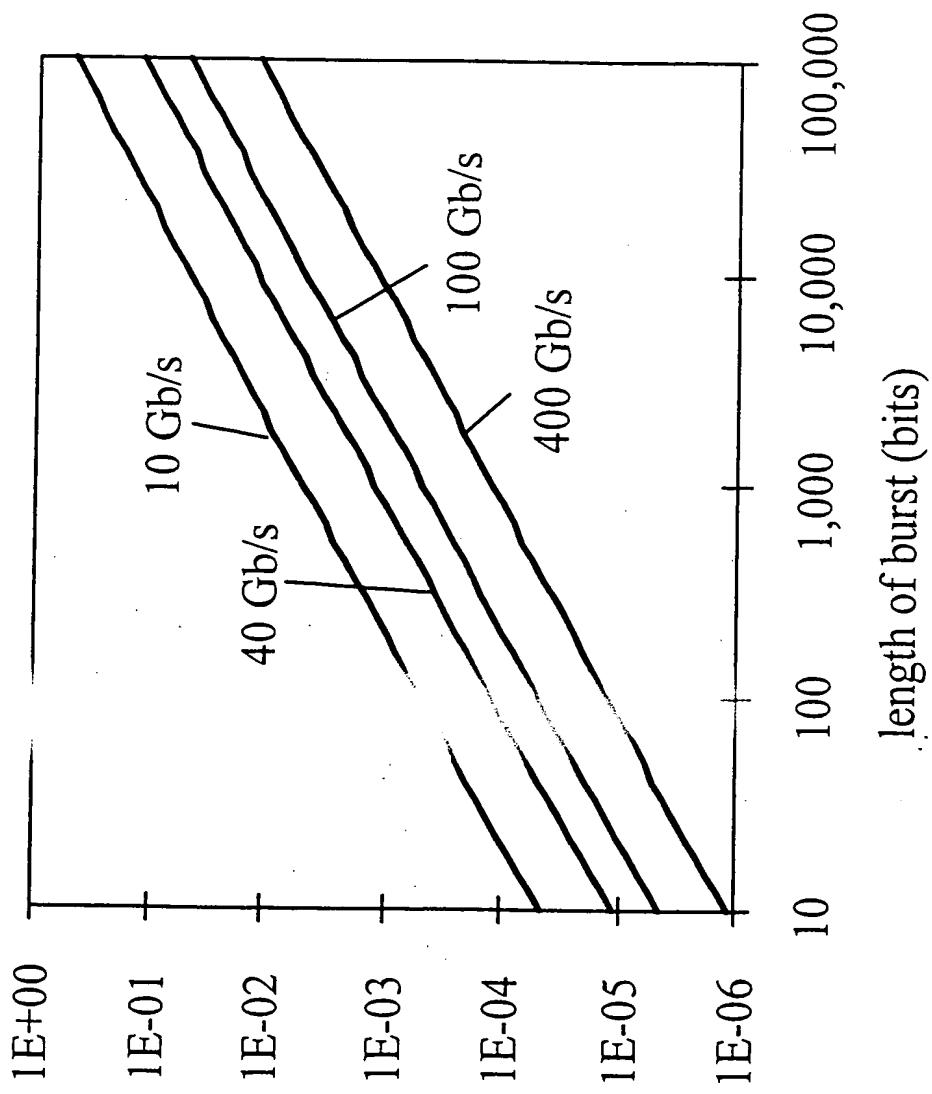


Figure 2

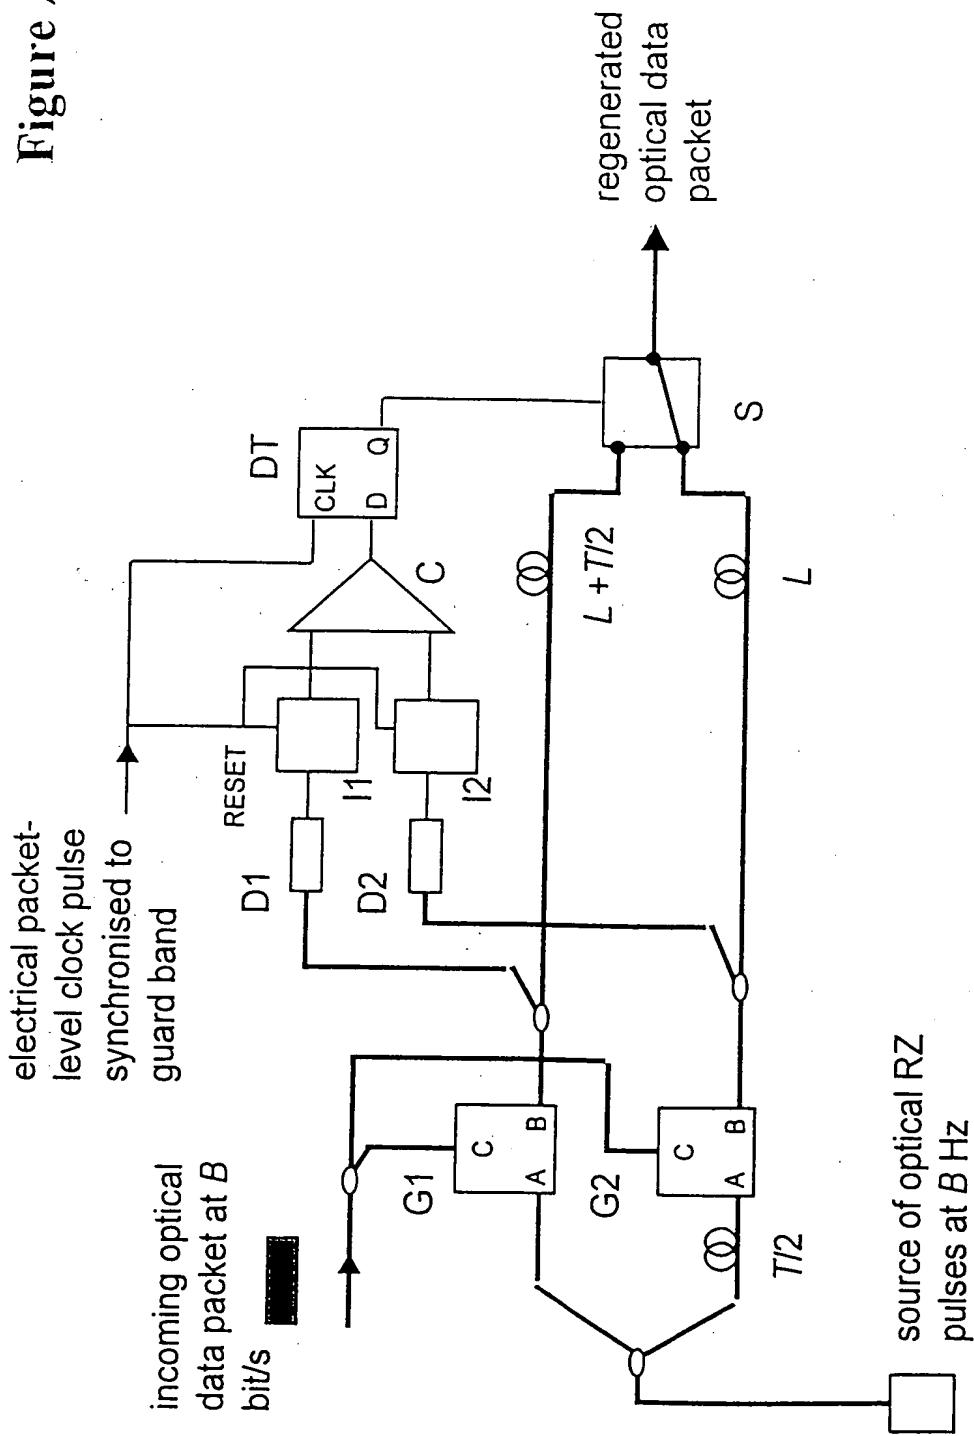
This Page Blank (uspto)

Figure 3



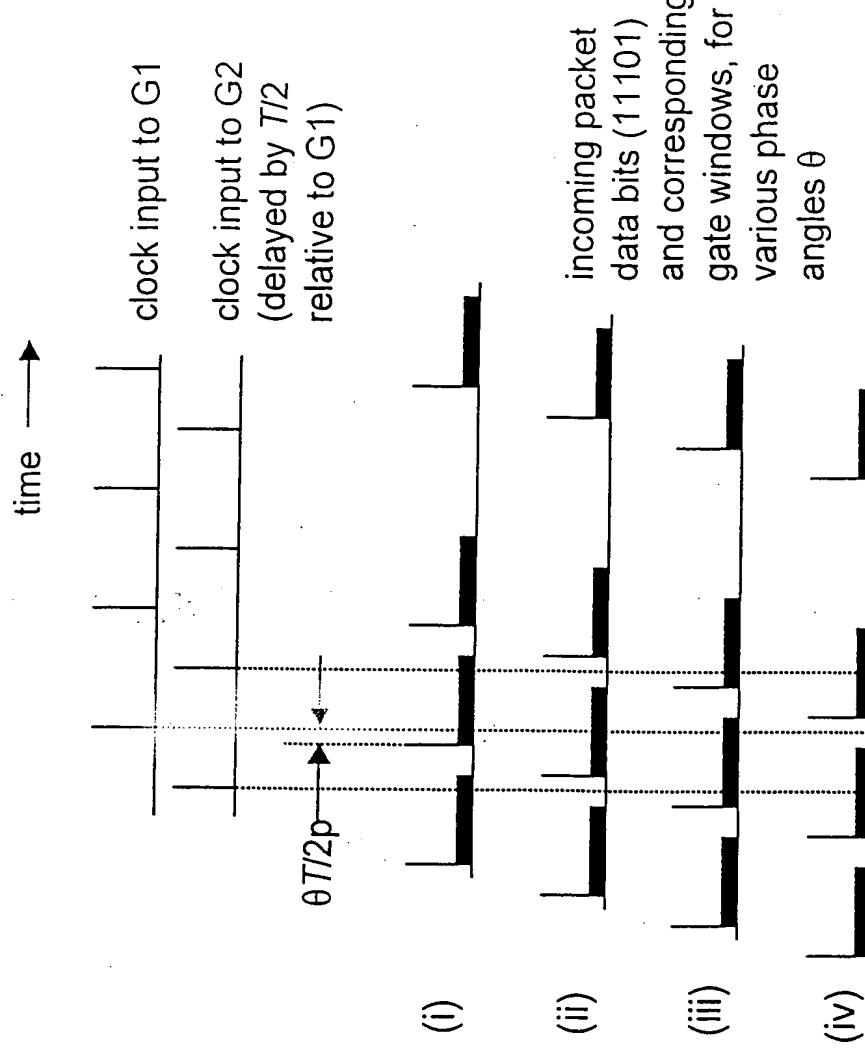
This Page Blank (uspto)

Figure 4



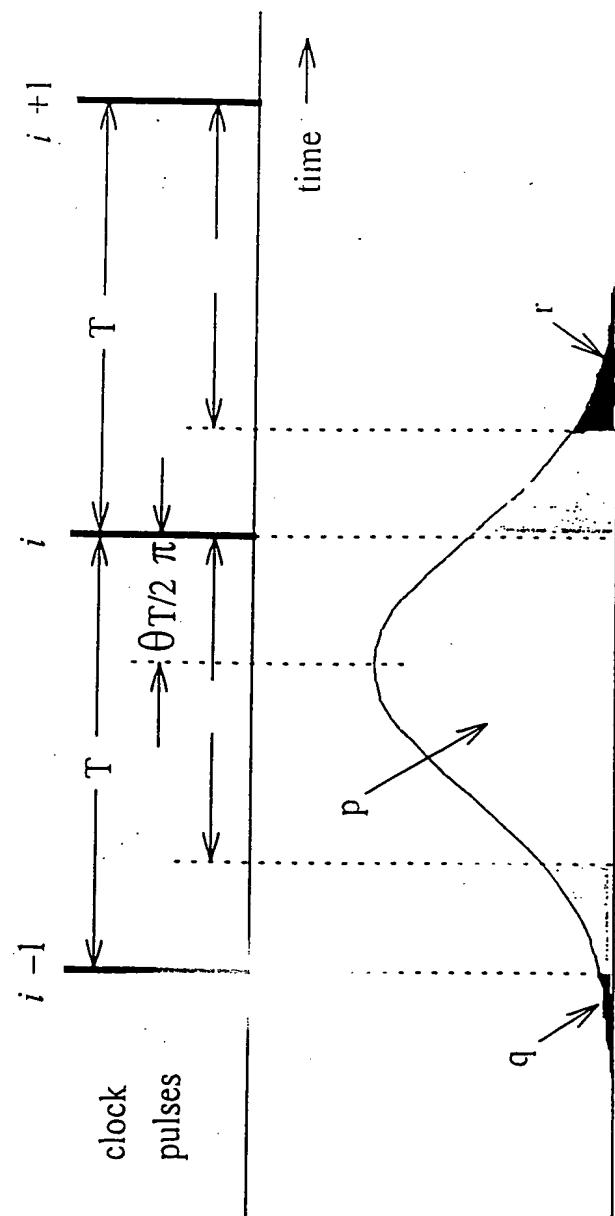
This Page Blank (uspto)

Figure 5



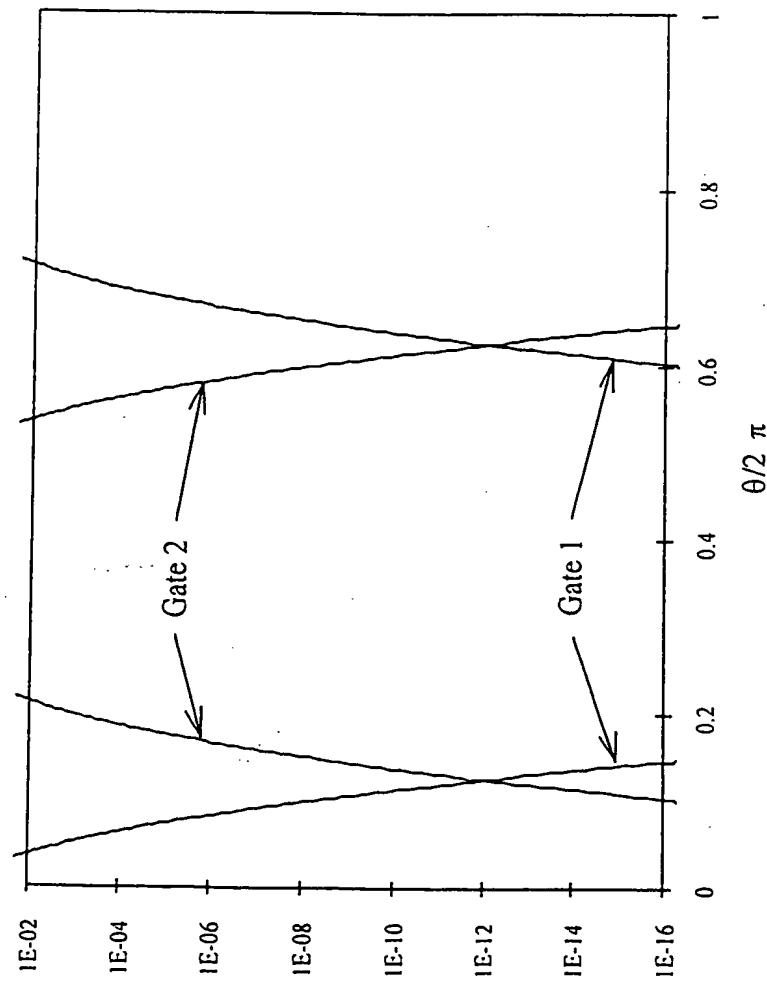
This Page Blank (uspto)

Figure φ



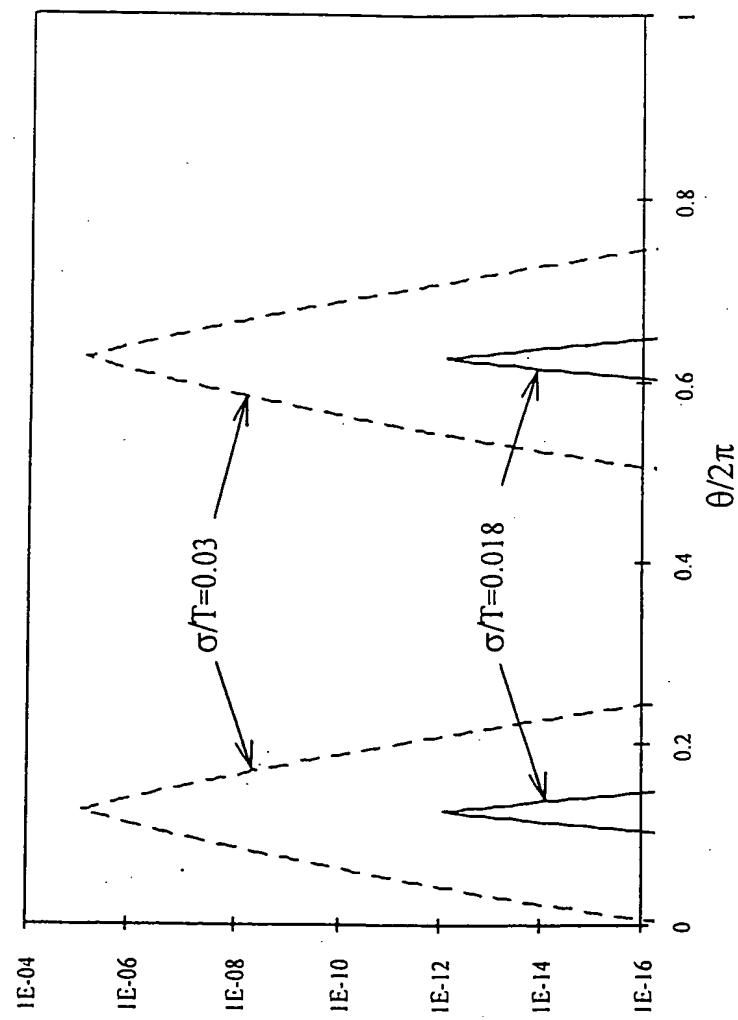
This Page Blank (uspto)

Figure 7



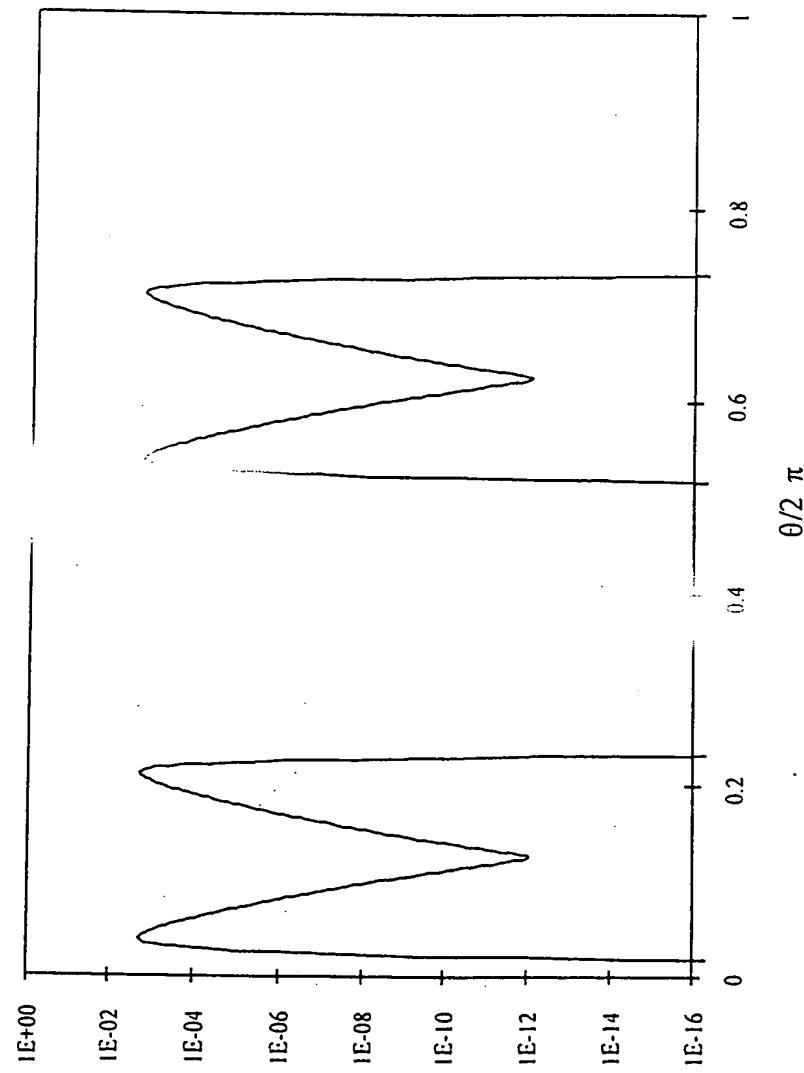
This Page Blank (uspto)

Figure 8



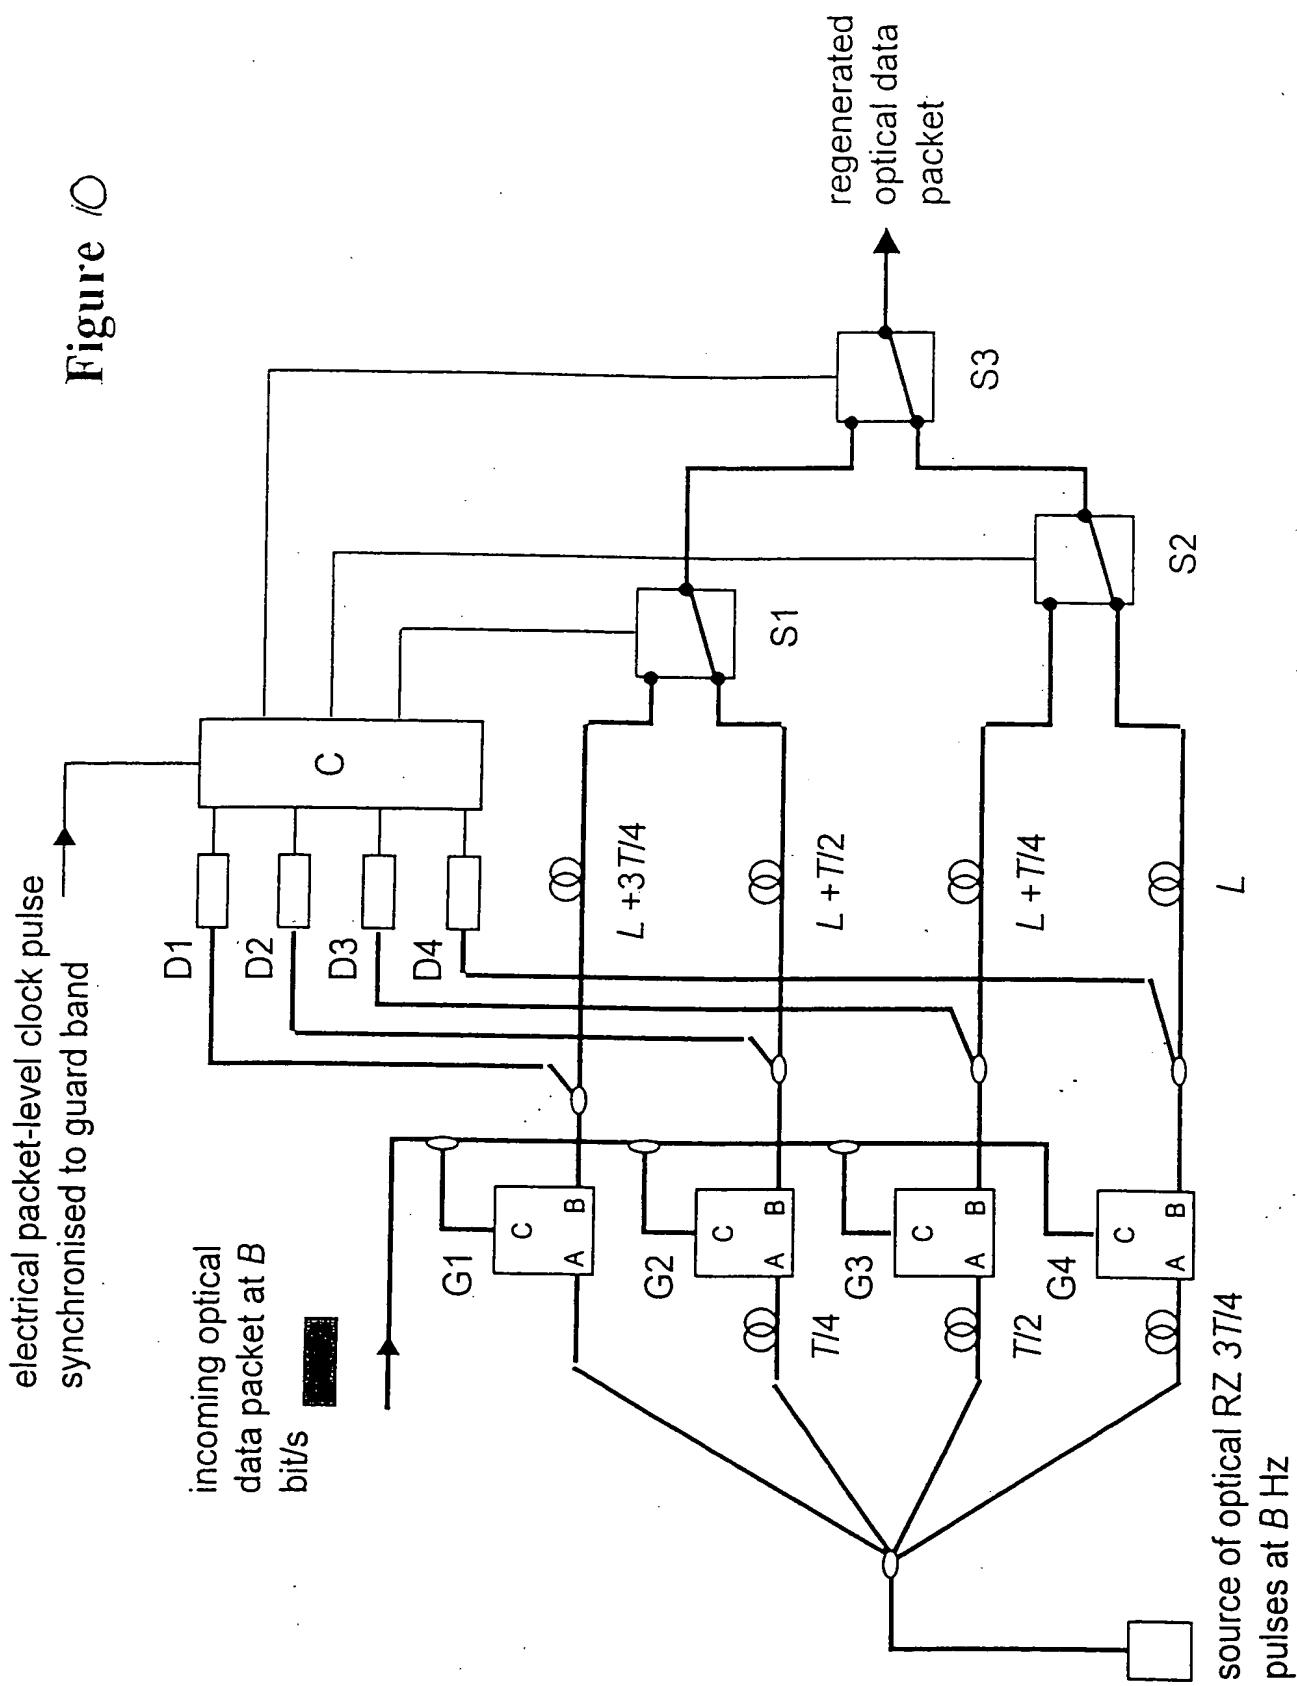
This Page Blank (uspto)

Figure 9



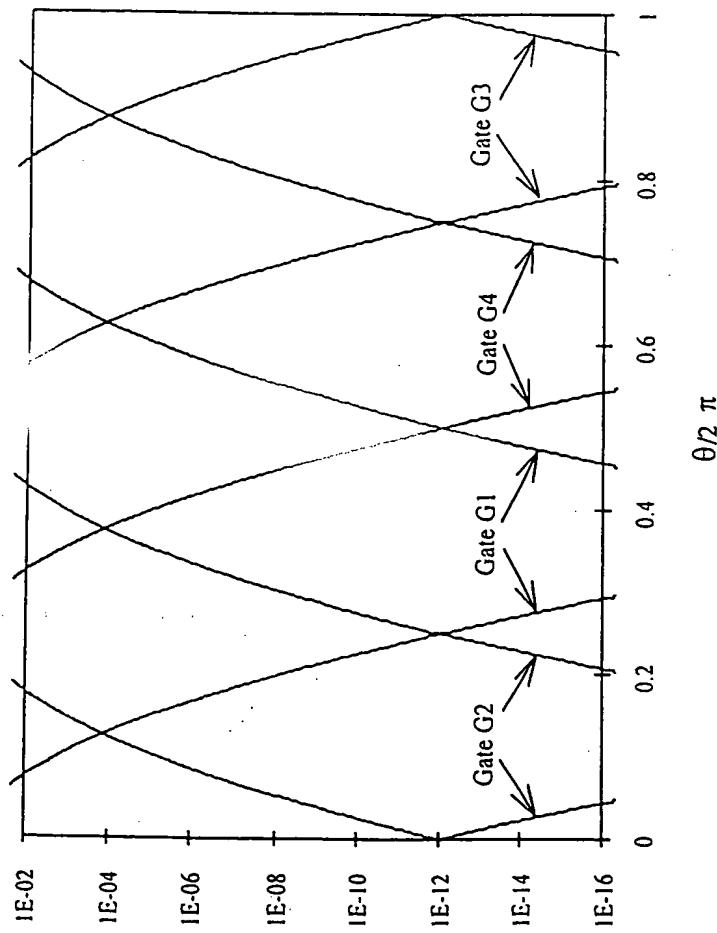
This Page Blank (uspto)

Figure 10



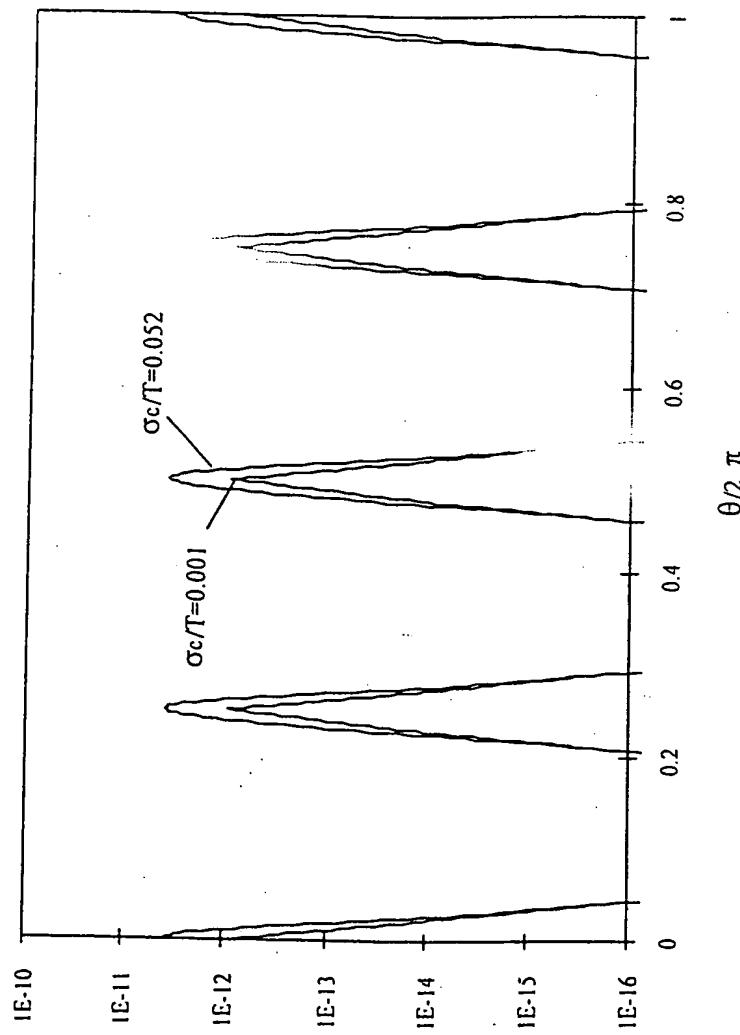
This Page Blank (uspto)

Figure 11



This Page Blank (uspto)

Figure 12



This Page Blank (uspto)

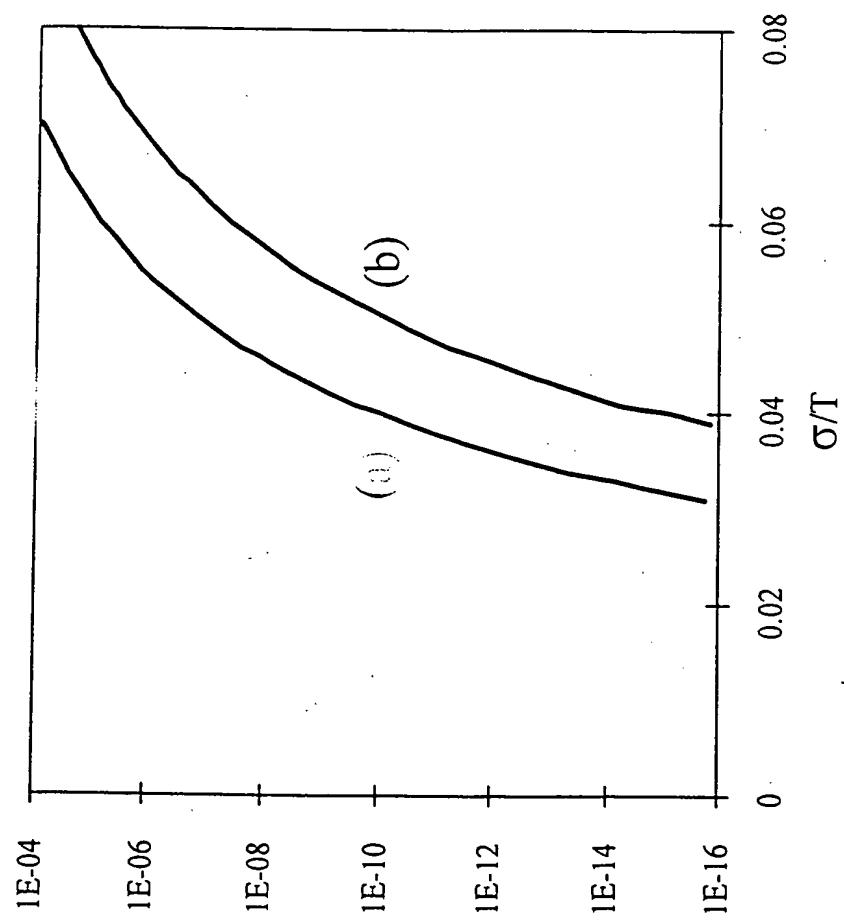
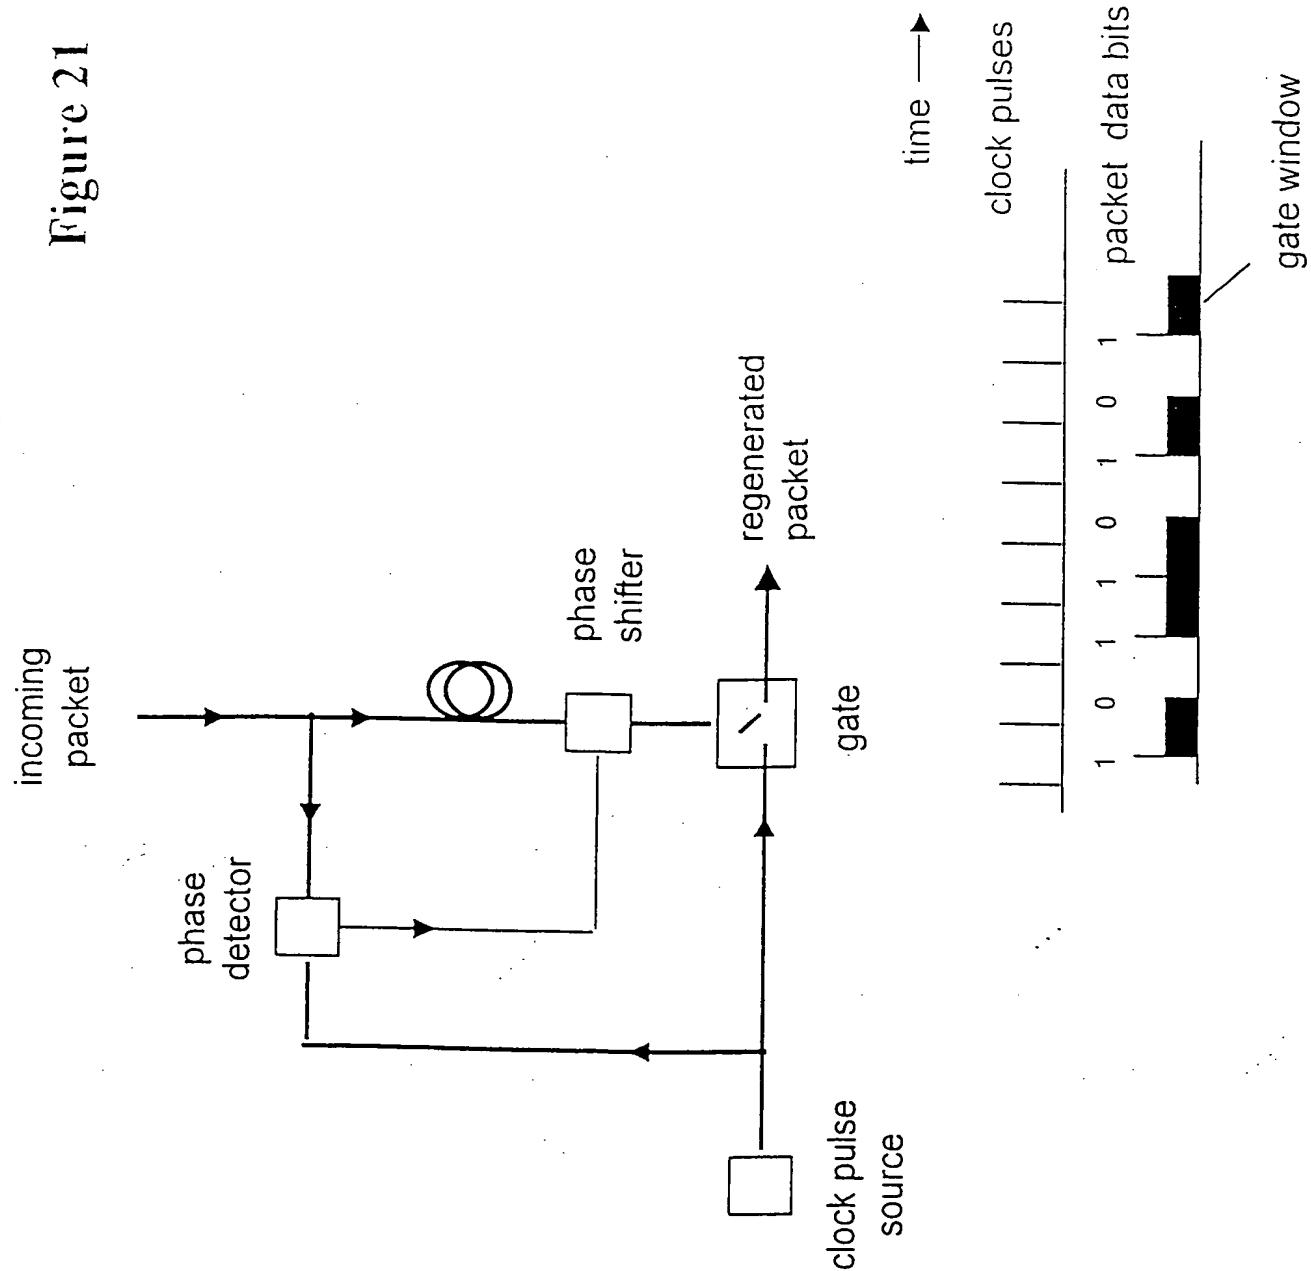


Figure 13

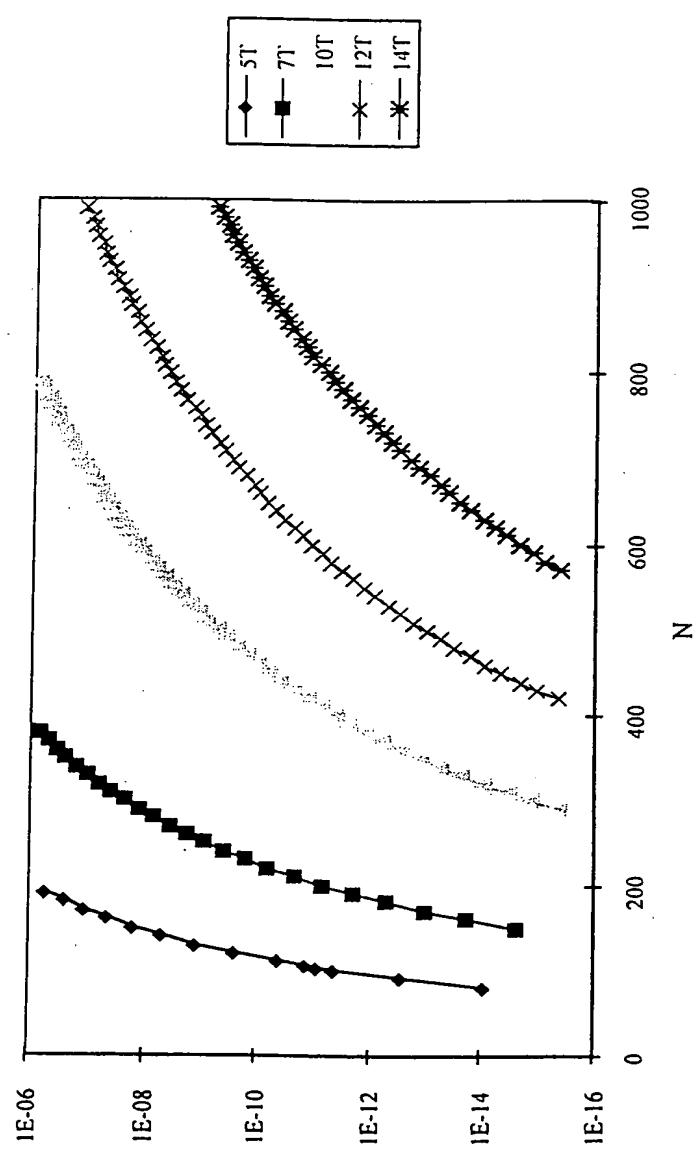
This Page Blank (uspto)

Figure 21



This Page Blank (uspto)

Figure 14



This Page Blank (uspto)

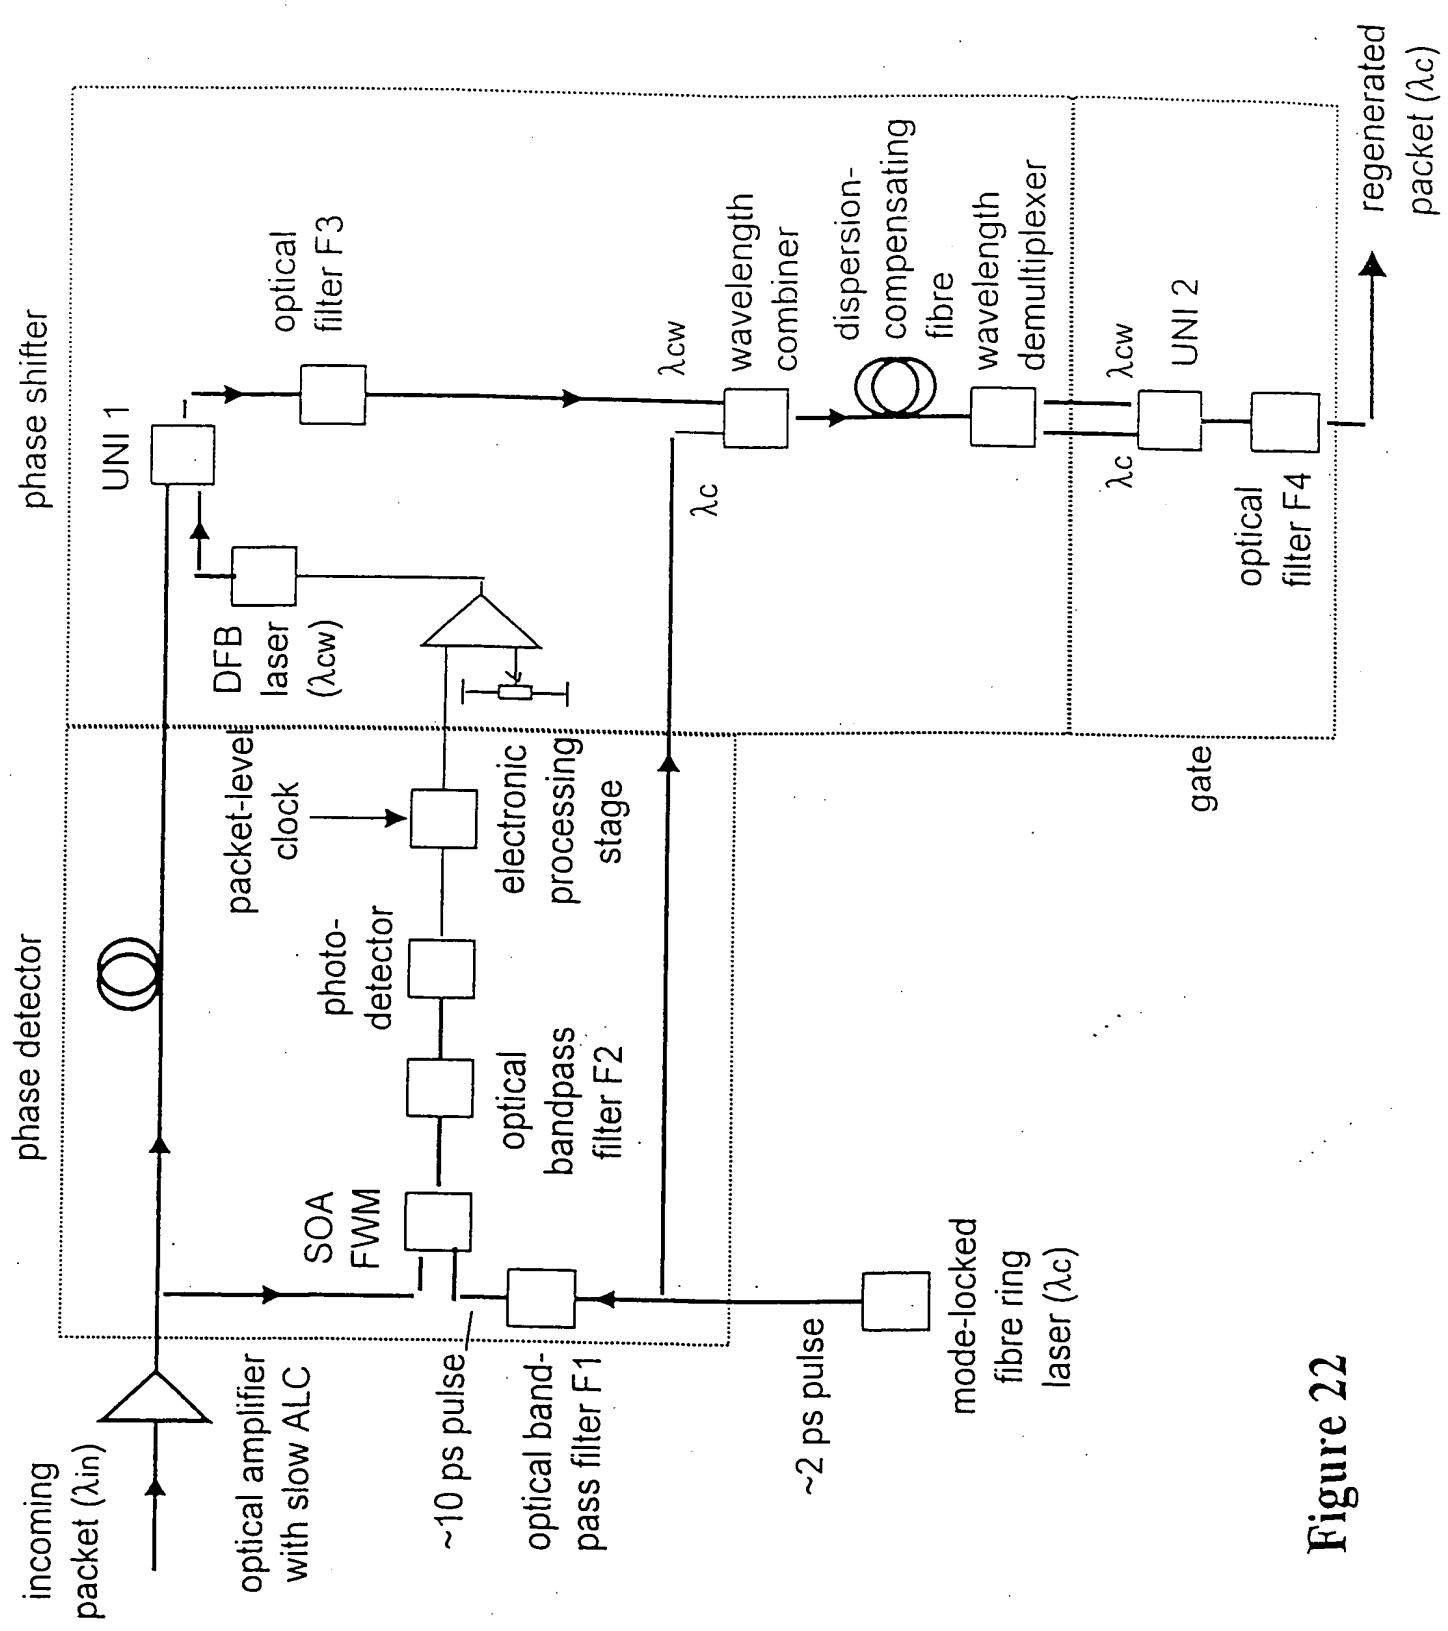


Figure 22

This Page Blank (uspto)

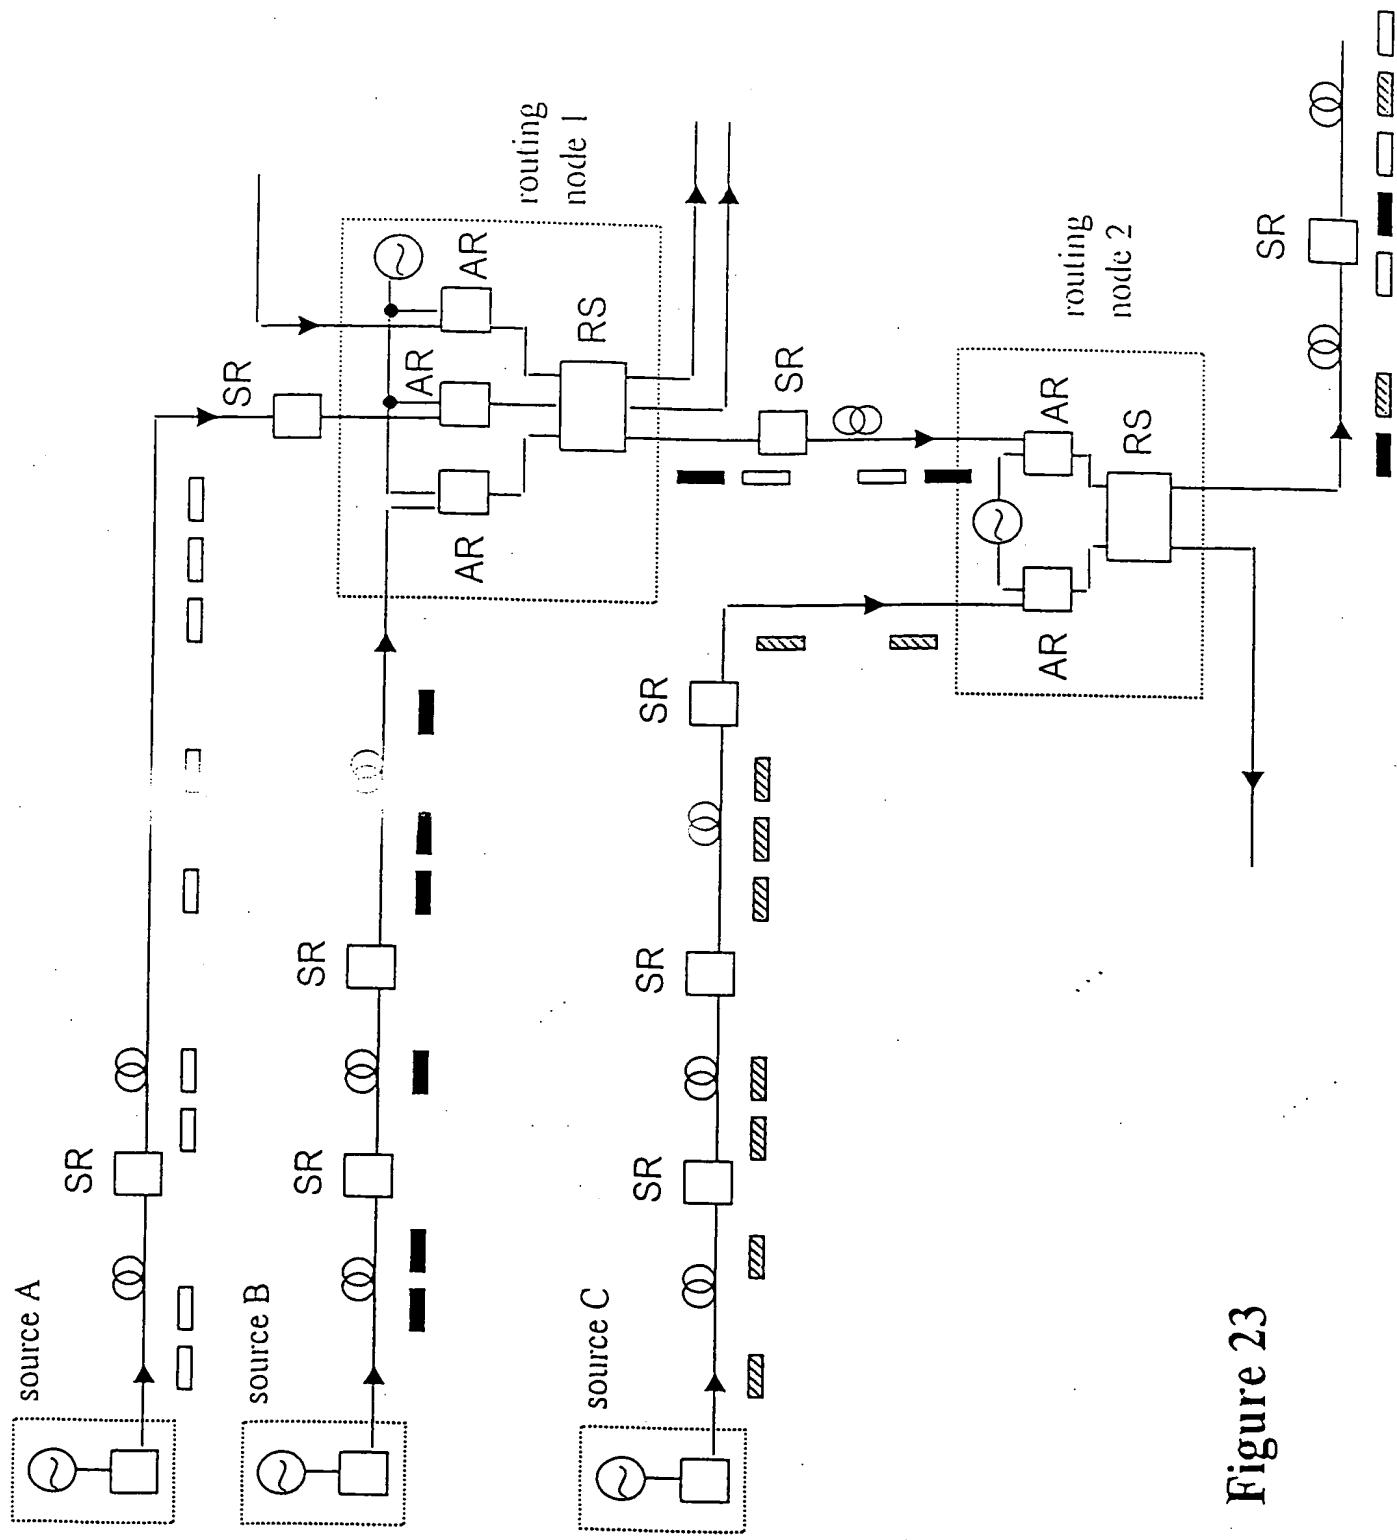
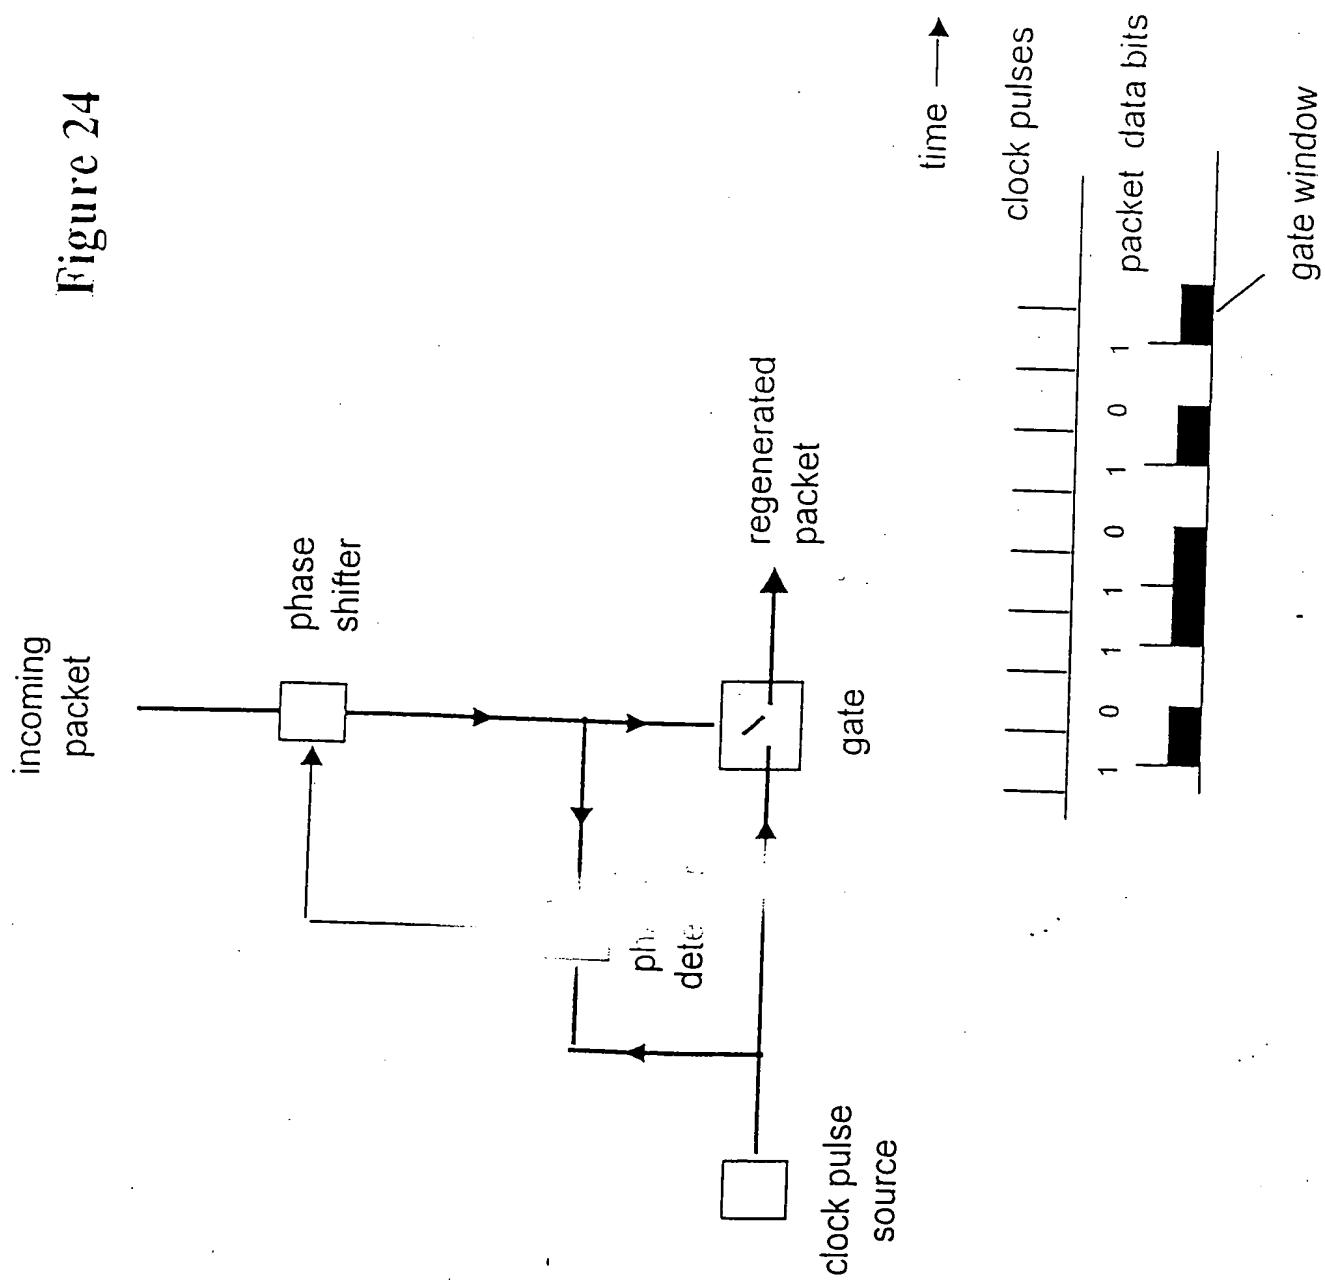


Figure 23

This Page Blank (uspto)

Figure 24



PC1 / GB 99/04159

15/4/99

BT Legal Group Services

This Page Blank (uspto)